arXiv:1404.0718v1 [cs.DM] 2 Apr 2014

Property Testing on Product Distributions: Optimal Testers for Bounded Derivative Properties

Deeparnab Chakrabarty* Kashyap Dixit† Madhav Jha‡ C. Seshadhri§

Abstract

The primary problem in property testing is to decide whether a given function satisfies a certain property, or is far from any function satisfying it. This crucially requires a notion of distance between functions. The most prevalent notion is the Hamming distance over the *uniform* distribution on the domain. This restriction to uniformity is more a matter of convenience than of necessity, and it is important to investigate distances induced by more general distributions. In this paper, we make significant strides in this direction. We give simple and optimal testers for *bounded derivative properties* over *arbitrary product distributions*. Bounded derivative properties include fundamental properties such as monotonicity and Lipschitz continuity. Our results subsume almost all known results (upper and lower bounds) on monotonicity and Lipschitz testing.We prove an intimate connection between bounded derivative property testing and binary search trees (BSTs). We exhibit a tester whose query complexity is the sum of expected depths of optimal BSTs for each marginal. Furthermore, we show this sum-of-depths is also a lower bound. A fundamental technical contribution of this work is an *optimal dimension reduction theorem* for all bounded derivative properties, which relates the distance of a function from the property to the distance of restrictions of the function to random lines. Such a theorem has been elusive even for monotonicity for the past 15 years, and our theorem is an exponential improvement to the previous best known result.

*Microsoft Research, dechakr@microsoft.com

†Pennsylvania State University, kashyap@cse.psu.edu, supported in part by NSF Grants CCF-0964655 and CCF-1320814

‡Sandia National Labs, Livermore, mjha@sandia.gov

§Sandia National Labs, Livermore, scomand@sandia.gov

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

1 Introduction

The field of property testing formalizes the following problem: how many queries are needed to decide if a given function satisfies a certain property? Formally, a *property* $\mathscr{P}$ is a subset of functions. A tester solves the relaxed membership problem of distinguishing functions in $\mathscr{P}$ from those 'far' from $\mathscr{P}$. To formalize 'far', one requires a notion of *distance*, $\mathsf{dist}(f, g) \in [0, 1]$, between functions. A function f is *ε-far from $\mathscr{P}$* if $\mathsf{dist}(f, g) \geq \varepsilon$ for all functions $g \in \mathscr{P}$. The notion of distance is central to property testing. The most prevalent notion of distance in the literature is the Hamming distance over the *uniform* distribution, that is, $\mathsf{dist}(f, g) := \Pr_{x \sim \mathscr{U}}[f(x) \neq g(x)]$, where $\mathscr{U}$ is the uniform distribution over the input domain. But the restriction to uniformity is more a matter of convenience than of necessity, and it is important and challenging to investigate distances induced by more general distributions. This was already underscored in the seminal work of Goldreich et. al. [GGR98] who "stressed that the generalization of property testing to arbitrary distributions" is important for applications. Nevertheless, a vast majority of results in property testing have focused solely on the uniform distribution. In this paper we investigate property testing of functions defined over the hypergrid $[n]^d$ with respect to distances induced by *arbitrary product distributions*. Product distributions over this domain form a natural subclass of general distributions where each individual coordinate is an arbitrary distribution independent of the other coordinates. They arise in many applications; the following are a couple of concrete ones. *Differential privacy:* Recent work on testing differential privacy [DJRT13] involve product distributions over the domain $[n]^d$. In this application, each domain point represents a database and each coordinate is a single individual's data. A distribution on databases is given by independent priors on each individual. The goal in [DJRT13] is to distinguish private mechanisms from those that aren't private on 'typical' databases. *Random testing of hardware:* Given an actual silicon implementation of a circuit, it is standard practice for engineers to test it on a set of random instances. Coordinates represent entities like memory addresses, data, control flow bits, etc. One chooses an independent but not identical distribution over each input to generate realistic set of test cases. There are specific commands in hardware languages like VHDL and Systemverilog [VHD02, ver14] that specify such coordinate-wise distributions. From a theoretical perspective, the study of property testing over non-uniform distributions has mostly led to work on specific problems. For uniform distributions, it is known that broad classes of algebraic and graphic properties are testable [AS08, AFNS09, KS08a, BSS10]. But little is known in this direction even for product distributions. One reason for this may be aesthetics: a priori, one doesn't expect a succinct, beautiful answer for testing over an arbitrary product distribution. In this paper we make significant strides in property testing under arbitrary product distributions. We give *simple*, *optimal* testers for the class of *bounded derivative* properties. This class contains (and is inspired by) the properties of monotonicity and Lipschitz continuity, which are of special interest in property testing. In fact, the *same* tester works for all such properties. Furthermore, our 'answer' is aesthetically pleasing: the optimal query complexity with respect to a product distribution is the sum of optimal binary search tree depths over the marginals. In particular, our results resolve a number of open problems in monotonicity testing, and subsume all previous upper and lower bounds over any product, including the uniform, distribution.

Previous Work. We set some context for our work. The property of monotonicity is simple. There is a natural coordinate-wise partial order over $[n]^d$. For a monotone function, $x \prec y$ im-

plies $f(x) \leq f(y)$. Monotonicity is one of the most well-studied properties in the area [EKK+00, GGL+00, DGL+99, LR01, FLN+02, AC06, Fis04, HK08b, PRR06, ACCL07, BRW05, BGJ+09, BCGSM12, BBM12, CS13a, CS13b, BRY14b]. A function is c-Lipschitz continuous if for all x, y, $|f(x) - f(y)| \leq c\|x - y\|_1$. Lipschitz continuity is a fundamental mathematical property with applications to differential privacy and program robustness. The study of Lipschitz continuity in property testing is more recent [JR11, AJMR12, CS13a, DJRT13, BRY14b]. With the exception of [HK07, HK08b, AC06, DJRT13], all the previous works are in the uniform distribution setting, for which the story is mostly clear: there is an $O(\varepsilon^{-1} d \log n)$-query tester for both properties [CS13a], and this is optimal for monotonicity [CS13b]. For Lipschitz continuity, an $\Omega(d \log n)$ non-adaptive lower bound has been proved [BRY14b] recently. For general product distributions, the story has been far less clear. Ailon and Chazelle [AC06] design an $O(2^d H/\varepsilon)$-query tester for monotonicity over product distributions, where H is the Shannon entropy of the distribution. This work connects property testing with information theory, and the authors explicitly ask whether the entropy is the "correct answer". There are no non-trivial lower bounds known for arbitrary product distributions. For Lipschitz continuity, no upper or lower bounds are known for general hypergrids, although an $O(d^2)$-query tester is known for the hypercube ($\{0,1\}^d$) domain [DJRT13]. Halevy and Kushilevitz [HK08b, HK07] study monotonicity testing in the *distribution-free setting*, where the tester does not know the input distribution but has access to random samples. Pertinent to us, they show a lower bound of $\Omega(2^d)$ for monotonicity testing over *arbitrary distributions*; for product distributions, Ailon and Chazelle [AC06] give an $O(\varepsilon^{-1} d 2^d \log n)$-query distribution-free tester.

1.1 Bounded Derivative Properties

To describe the class of bounded derivative properties, we first set some notation. For an integer k, we use $[k]$ to denote the set $\{1, 2, \ldots, k\}$. Consider a function $f : [n]^d \mapsto \mathbb{R}$ and a dimension $r \in [d]$. Define $\partial_r f(x) := f(x + \mathbf{e}_r) - f(x)$, where $\mathbf{e}_r$ is the unit vector in the rth dimension ($\partial_r f$ is defined only on x with $x_r < n$.).

Definition 1.1. *An ordered set* $\mathbf{B}$ *of* $2d$ *functions* $l_1, u_1, l_2, u_2, \ldots, l_d, u_d : [n-1] \mapsto \mathbb{R}$ *is called a* bounding family *if for all* $r \in [d]$ *and* $y \in [n-1]$, $l_r(y) < u_r(y)$. *Let* $\mathbf{B}$ *be a bounding family of functions. The property of being* $\mathbf{B}$-derivative bounded*, denoted as* $\mathscr{P}(\mathbf{B})$*, is the set of functions* $f : [n]^d \mapsto \mathbb{R}$ *such that: for all* $r \in [d]$ *and* $x \in [n]^d$*,*

$$l_r(x_r) \leq \partial_r f(x) \leq u_r(x_r). \tag{1}$$

This means the rth-partial derivative of f is bounded by quantities that only depend on the rth coordinate. Note that this dependence is completely arbitrary, and different dimensions can have completely different bounds. This forms a rich class of properties which includes monotonicity and c-Lipschitz continuity. To get monotonicity, simply set $l_r(y) = 0$ and $u_r(y) = \infty$ for all r. To get c-Lipschitz continuity, set $l_r(y) = -c$ and $u_r(y) = +c$ for all r. The class also includes the property demanding monotonicity for some (fixed) coordinates and the c-Lipschitz continuity for others; and the non-uniform Lipschitz property that demands different Lipschitz constants for different coordinates.

Definition 1.2. *Fix a bounding family* $\mathbf{B}$ *and product distribution* $\mathscr{D} = \prod_{r \leq d} \mathscr{D}_r$*. Define* $\mathsf{dist}_{\mathscr{D}}(f, g) = \Pr_{x \sim \mathscr{D}}[f(x) \neq g(x)]$*. A property tester for* $\mathscr{P}(\mathbf{B})$ *with respect to* $\mathscr{D}$ *takes as input proximity parameter* $\varepsilon > 0$ *and has query access to function* f*. If* $f \in \mathscr{P}(\mathbf{B})$*, the tester accepts with probability* $> 2/3$*. If* $\mathsf{dist}_{\mathscr{D}}(f, \mathscr{P}(\mathbf{B})) > \varepsilon$*, the tester rejects with probability* $> 2/3$*.*

1.2 Main Results

Our primary result is a property tester for all bounded-derivative properties over any product distribution. The formal theorem requires some definitions of search trees. Consider any binary search tree (BST) T over the universe $[n]$, and let the depth of a node denote the number of *edges* from it to the root. For a distribution $\mathscr{D}_r$ over $[n]$, the *optimal BST for* $\mathscr{D}_r$ is the BST minimizing the expected depth of vertices drawn from $\mathscr{D}_r$. Let $\Delta^*(\mathscr{D}_r)$ be this optimal depth: a classic dynamic programming solution finds this optimal tree [Knu73, Yao82] in polynomial time. Given a product distribution $\mathscr{D} = \prod_{r \le d} \mathscr{D}_r$, we abuse notation and let $\Delta^*(\mathscr{D})$ denote the sum $\sum_{r=1}^{d} \Delta^*(\mathscr{D}_r)$.

Theorem 1.3. [Main upper bound] *Consider functions $f : [n]^d \mapsto \mathbb{R}$. Let $\mathbf{B}$ be a bounding family and $\mathscr{D}$ be a product distribution. There is a tester for $\mathscr{P}(\mathbf{B})$ w.r.t. $\mathscr{D}$ making $100\varepsilon^{-1}\Delta^*(\mathscr{D})$ queries.*

The tester is non-adaptive with one-sided error, that is, the queries don't depend on the answers, and the tester always accepts functions satisfying the property. Furthermore, the *same* tester works for all bounding families, that is, the set of queries made by the tester doesn't depend on $\mathbf{B}$. Interestingly, the "worst" distribution is the uniform distribution, where $\Delta^*(\mathscr{D})$ is maximized to $\Theta(d \log n)$. We remark that the class of bounded derivative properties was not known to be testable even under uniform distributions. Results were known [CS13a] (only under the uniform distribution) for the subclass where all l_r (and u_r) are the same, constant function. To give perspective on the above result, it is instructive to focus on say just monotonicity (one can repeat this for Lipschitz). Let $H(\mathscr{D})$ denote the Shannon entropy of distribution $\mathscr{D}$ over the hypergrid. It is well-known that $\Delta^*(\mathscr{D}_r) \le H(\mathscr{D}_r)$ (see [Meh75] for a proof), so $\Delta^*(\mathscr{D}) \le H(\mathscr{D})$ for product distribution $\mathscr{D}$.

Corollary 1.4. *Consider functions $f : [n]^d \mapsto \mathbb{R}$. Monotonicity testing over a product distribution $\mathscr{D}$ can be done with $100H(\mathscr{D})/\varepsilon$ queries.*

This is an *exponential* improvement over the previous best result of Ailon and Chazelle [AC06], who give a monotonicity tester with query complexity $O(2^d H(\mathscr{D})/\varepsilon)$. Observe that for uniform distributions, $H(\mathscr{D}) = \Theta(d \log n)$, and therefore the above result subsumes the optimal testers of [CS13a]. Now consider the monotonicity testing over the boolean hypercube.

Corollary 1.5. *Consider functions $f : \{0,1\}^d \mapsto \mathbb{R}$. Monotonicity testing over any product distribution $\mathscr{D} = \prod_{r=1}^{d} \mathscr{D}_r$, where each $\mathscr{D}_r = (\mu_r, 1-\mu_r)$, can be done with $100\varepsilon^{-1} \sum_{r=1}^{d} \min(\mu_r, 1-\mu_r)$ queries.*

Given that monotonicity testing over the hypercube has received much attention [GGL+00, DGL+99, LR01, FLN+02, BBM12, CS13a, CS13b], it is somewhat surprising that *nothing non-trivial* was known even over the p-biased distribution for $p \neq 1/2$; our result implies an $O(\varepsilon^{-1}pd)$-query tester. The above corollary also asserts that entropy of a distribution *doesn't capture the complexity of monotonicity testing* since the entropy, $\sum_r \mu_r \log(1/\mu_r) + (1-\mu_r)\log(1/(1-\mu_r))$, can be larger than the query complexity described above by a logarithmic factor. For example, if each $\mu_r = 1/\sqrt{d}$, the tester of Corollary 1.5 requires $O(\sqrt{d}/\varepsilon)$ queries, while $H(\mathscr{D}) = \Theta(\sqrt{d}\log d)$. We complement Theorem 1.3 with a matching lower bound, cementing the connection between testing of bounded-derivative properties and optimal search tree depths. This requires a technical definition of stable distributions, which is necessary for the lower bound. To see this consider

a distribution $\mathscr{D}$ for which there exists a product distribution $\mathscr{D}'$ such that $||\mathscr{D}' - \mathscr{D}||_{\mathsf{TV}} \leq \varepsilon/2$ but $\Delta^*(\mathscr{D}') \ll \Delta^*(\mathscr{D})$. One could simply apply Theorem 1.3 with $\mathscr{D}'$ to obtain a tester with a much better query complexity than $\Delta^*(\mathscr{D})$. $\mathscr{D}$ is called (ε', ρ)-stable if $\|\mathscr{D} - \mathscr{D}\| \leq \varepsilon'$ implies $\Delta^*(\mathscr{D}') \geq \rho\Delta^*(\mathscr{D})$, for any product distribution $\mathscr{D}'$.

Theorem 1.6. [Main lower bound] *For any parameter ε, there exists $\varepsilon' = \Theta(\varepsilon)$ such that for any bounding family $\mathbf{B}$ and (ε', ρ)-stable, product distribution $\mathscr{D}$, any (even adaptive, two-sided) tester for $\mathscr{P}(\mathbf{B})$ w.r.t. $\mathscr{D}$ with proximity parameter ε requires $\Omega(\rho\Delta^*(\mathscr{D}))$ queries.*

This lower bound is new even for monotonicity testing over one dimension. Ailon and Chazelle [AC06] explicitly ask for lower bounds for monotonicity testing for domain $[n]$ over arbitrary distributions. Our upper and lower bounds completely resolve this problem. For Lipschitz testing, the state of the art was a *non-adaptive* lower bound of $\Omega(d \log n)$ for the uniform distribution [BRY14b]. Since the uniform distribution is stable, the previous theorem implies an optimal $\Omega(d \log n)$ lower bound even for adaptive, two-sided testers over the uniform distribution. The previous upper bounds are in the setting where the tester knows the distribution $\mathscr{D}$. In the *distribution-free* setting, the tester only gets random samples from $\mathscr{D}$ although it is free to query any point of the domain. As a byproduct of our approach, we also get results for this setting. The previous best bound was an $O(\varepsilon^{-1} d 2^d \log n)$ query tester [AC06].

Theorem 1.7. *Consider functions $f : [n]^d \mapsto \mathbb{R}$. There is a distribution-free (non-adaptive, one-sided) tester for $\mathscr{P}(\mathbf{B})$ w.r.t. $\mathscr{D}$ making $100\varepsilon^{-1} d \log n$ queries.*

1.3 Technical highlights

Optimal dimension reduction. The main engine running the upper bounds is an optimal dimension reduction theorem. Focus on just the uniform distribution. Given $f : [n]^d \mapsto \mathbb{R}$ that is ε-far from $\mathscr{P}(\mathbf{B})$, what is the expected distance of the function restricted to a uniform random line in $[n]^d$? This natural combinatorial question has been at the heart of various monotonicity testing results [GGL+00, DGL+99, AC06, HK08b]. The best known bounds are that this expected distance is at least $\varepsilon/(d2^d)$ [AC06, HK08b]. Weaker results are known for the Lipschitz property [JR11, AJMR12]. We given an optimal resolution (up to constant factors) to this problem not only for the uniform distribution, but for any arbitrary product distribution, and for any bounded derivative property. In $[n]^d$, an r-line is a combinatorial line parallel to the r-axis. Fix some bounding family $\mathbf{B}$ and product distribution $\mathscr{D} = \prod_r \mathscr{D}_r$. Note that $\mathscr{D}_{-r} = \prod_{i \neq r} \mathscr{D}_i$ is a distribution on r-lines. If we restrict f to an r-line ℓ, we get a function $f|_\ell : [n] \mapsto \mathbb{R}$. It is meaningful to look at the distance of $f|_\ell$ to $\mathscr{P}(\mathbf{B})$ (though this only involves the bounds of $l_r, u_r \in \mathbf{B}$). Let $\mathsf{dist}^r_{\mathscr{D}}(f, \mathscr{P}(\mathbf{B})) := \mathbf{E}_{\ell \sim \mathscr{D}_{-r}}[\mathsf{dist}_{\mathscr{D}_r}(f|_\ell, \mathscr{P}(\mathbf{B}))]$.

Theorem 1.8. [Optimal Dimension Reduction] *Fix bounding family $\mathbf{B}$ and product distribution $\mathscr{D}$. For any function f,*

$$\sum_{r=1}^{d} \mathsf{dist}^r_{\mathscr{D}}(f, \mathscr{P}(\mathbf{B})) \geq \mathsf{dist}_{\mathscr{D}}(f, \mathscr{P}(\mathbf{B}))/4.$$

Let us give a short synopsis of previous methods used to tackle the case of monotonicity in the uniform distribution case. For brevity's sake, let ε^r_f denote $\mathsf{dist}^r_{\mathscr{U}}(f, \texttt{MON})$ and ε_f denote

$\mathsf{dist}_{\mathscr{U}}(f, \mathtt{MON})$. That is, $\varepsilon_f^r n^d$ modifications makes the function monotone along the r-dimension, and the theorem above states that $4\sum_r \varepsilon_f^r n^d$ modifications suffice to make the whole function monotone. Either explicitly or implicitly, previous attempts have taken a constructive approach: they use the modifications along the rth dimensions to correct the whole function. Although in principle a good idea, a bottleneck to the above approach is that correcting the function along one dimension may potentially introduce significantly larger errors along other dimensions. Thus, one can't just "add up" the corrections in a naive manner. The process is even more daunting when one tries this approach for the Lipschitz property. Our approach is completely different, and is 'non-constructive', and looks at all bounded-derivative properties in a uniform manner. We begin by proving Theorem 3.1 for $\mathscr{P}(\mathbf{B})$ over the uniform distribution. The starting point is to consider a weighted violation graph G, where any two domains point forming a violation to $\mathscr{P}(\mathbf{B})$ are connected (the weight is a "magnitude" of violation). It is well-known that the size of a maximum matching M in G is at least $\varepsilon_f n^d/2$. The main insight is to use different matchings to get handles on the distance ε_f^r rather than using modifications that correct the function. More precisely, we construct a sequence of special matchings $M = M_0, M_1, \ldots, M_d = \emptyset$, such that the drop in size $|M_{r-1}| - |M_r|$ is at most $2\varepsilon_f^r n^d$, which proves the above theorem. This requires structural properties on the M_r's proven using the *alternating path machinery* developed in [CS13a]. What about a general product distribution $\mathscr{D}$? Suppose we 'stretch' every point in every direction proportional to its marginal. This leads to a 'bloated' hypergrid $[N]^d$ where each point in the original hypergrid corresponds to a high-dimensional cuboid. By the obvious association of function values, one obtains a $f_{\mathtt{ext}} : [N]^d \mapsto \mathbb{R}$. If $\mathscr{P}(\mathbf{B})$ is monotonicity, then it is not hard to show that $\mathsf{dist}_{\mathscr{D}}(f) = \mathsf{dist}_{\mathscr{U}}(f_{\mathtt{ext}})$. So we can apply dimension reduction for $f_{\mathtt{ext}}$ over the uniform distribution and map it back to f over $\mathscr{D}$. However, such an argument breaks down for Lipschitz (let alone general $\mathbf{B}$), since $\mathsf{dist}_{\mathscr{U}}(f')$ can be much smaller than $\mathsf{dist}_{\mathscr{D}}(f)$. The optimal fix for $f_{\mathtt{ext}}$ could perform non-trivial changes within the cuboidal regions, and this cannot be mapped back to a fix for the original f. This is where the generality of the bounded-derivative properties saves the day. For any $\mathbf{B}$ and $\mathscr{D}$, we can define a new bounding family $\mathbf{B}_{\mathtt{ext}}$ over $[N]^d$, such that $\mathsf{dist}_{\mathscr{D}}(f, \mathscr{P}(\mathbf{B})) = \mathsf{dist}_{\mathscr{U}}(f_{\mathtt{ext}}, \mathscr{P}(\mathbf{B}_{\mathtt{ext}}))$. Now, dimension reduction is applied to $f_{\mathtt{ext}}$ for $\mathscr{P}(\mathbf{B}_{\mathtt{ext}})$ over $\mathscr{U}$ and translated back to the original setting. **Search trees and monotonicity.** An appealing aspect of our results is the tight connection between optimal search trees over product distributions to bounded-derivative properties. The dimension reduction lemma allows us (for the upper bounds) to focus on just the line domain $[n]$. For monotonicity testing on $[n]$ over an arbitrary distribution $\mathscr{D}$, Halevy and Kushilevitz gave an $O(\varepsilon^{-1} \log n)$-query distribution free tester [HK08b], and Ailon and Chazelle gave an $O(\varepsilon^{-1} H(\mathscr{D}))$-query tester [AC06]. Pretty much every single result for monotonicity testing on $[n]$ involves some analogue of binary search [EKK+00, BRW05, ACCL07, PRR06, HK08b, AC06, BGJ+09]. But we make this connection extremely precise. We show that *any* binary search tree can be used to get a tester with respect to an arbitrary distribution, whose expected query complexity is the expected depth of the tree with respect to the distribution. This argument is extremely simple in hindsight, but it is a significant conceptual insight. Firstly, it greatly simplifies earlier results – using the completely balanced BST, we get an $O(\varepsilon^{-1} \log n)$-distribution free tester; with the optimal BST, we get $O(\varepsilon^{-1} H(\mathscr{D}))$-queries. The BST tester along with the dimension reduction, provides a tester for $[n]^d$ whose running time can be better than $H(\mathscr{D})$ (especially for the hypercube). Most importantly, optimal BSTs are a crucial ingredient for our lower bound construction. **Lower Bounds for Product Distributions.** The first step to general lower bounds is a simple reduction from monotonicity testing to any bounded-derivative property. Again, the reduction may seem trivial in

hindsight, but note that special sophisticated constructions were used for existing Lipschitz lower bounds [JR11, BRY14b]. For monotonicity, we use the framework developed in [Fis04, CS13b] that allows us to focus on comparison based testers. The lower bound for $[n]$ uses a convenient near-optimal BST. For each level of this tree we construct a 'hard' non-monotone function, leading to (roughly) $\Delta^*(\mathscr{D})$ such functions in case of stable distributions. These functions have violations to monotonicity lying in 'different regions' of the line, and any bonafide tester must make a different query to catch each function. In going to higher dimensions, we face a significant technical hurdle. The line lower bound easily generalizes to the hypergrid *if each marginal distribution is individually stable.* However, this may not be the case – there are stable product distributions whose marginals are unstable. As a result, each dimension may give 'hard' functions with very small distance. Our main technical contribution is to show how to *aggregate* functions from various dimensions together to obtain hard functions for the hypergrid in such a way that the distances add up. This is rather delicate, and is perhaps the most technical portion of this paper. In summary, we show that for stable distributions, the total search-tree depth is indeed the lower bound for testing monotonicity, and via the reduction mentioned above, for any bounded-derivative property.

1.4 Other Related Work.

Monotonicity testing has a long history, and we merely point the reader to the discussions in [CS13a, CS13b]. The work on testing over non-uniform distributions was performed in [HK08b, HK07, AC06], the details of which have been provided in the previous section. Goldreich et al. [GGR98] had already posed the question of testing properties of functions over non-uniform distributions, and obtain some results for dense graph properties. A serious study of the role of distributions was undertaken by Halevy and Kushilevitz [HK07, HK08b, HK05, HK08a], who formalized the concept of distribution-free testing. (Refer to Halevy's thesis [Hal06] for a comprehensive study.) Kopparty and Saraf extend the classic linearity test to classes of distributions, including product distributions [KS08b]. Glasner and Servedio [GS09] and Dolev and Ron [DR11] give various upper and lower bounds for distribution-free testers for various properties over $\{0,1\}^n$. Non-uniform distributions were also considered recently in the works of Balcan et al. [BBBY12] and [GR13] which constrain the queries that can be made by the tester to samples drawn from the distribution. Recent work of Berman et al. [BRY14a] introduces property testing over ℓ_p-distances. We believe work along these lines studying richer notions of distance is critical to the growth of property testing.

Note to the reader. The paper is rather long, although, we hope the extended introduction above will allow the reader to choose the order in which to peruse the paper. We give a brief outline of remainder. In §2, we define a particular *quasi-metric* corresponding to a bounding family $\mathbf{B}$ and give an equivalent definition of the bounded-derivative property with respect to it. This definition is convenient and will be the one used for the rest of the paper. This section must be read next. The dimension reduction theorem is presented in its full glory in §3. In §4.1, we describe the tester when the domain is just the line, and the easy generalization to the hypergrid via dimension reduction is presented in §4.2. For lower bounds, we prove the reduction to monotonicity in §5.1, and describe the approach to montonicity lower bounds in §5.2. The hard families for the line is given in §5.3, for the hypercube in §5.4, and the general hypergrid lowerbound is described in §5.5.

2 Quasimetric induced by a Bounding Family

It is convenient to abstract out $\mathscr{P}(\mathbf{B})$ in terms of a *metric-bounded property*. Such ideas was used in [CS13a] to give a unified proof for monotonicity and Lipschitz for the uniform distribution. The treatment here is much more general. We define a quasimetric depending on $\mathbf{B}$ denoted by $\mathfrak{m}(x, y)$.

Definition 2.1. *Given bounding family* $\mathbf{B}$*, construct the weighted directed hypergrid* $[n]^d$*, where all adjacent pairs are connected by two edges in opposite directions. The weight of* $(x + \mathbf{e}_r, x)$ *is* $u_r(x_r)$ *and the weight of* $(x, x + \mathbf{e}_r)$ *is* $-l_r(x_r)$*.* $\mathfrak{m}(x, y)$ *is the shortest path weight from* x *to* y*.*

Note that $\mathfrak{m}$ is asymmetric, can take negative values, and $\mathfrak{m}(x, y) = 0$ does not necessarily imply $x = y$. For these reasons, it is really a possibly-negative-pseudo-quasi-metric, although we will refer to it simply as a metric in the remainder of the paper. Since $\mathbf{B}$ is a bounding family, any cycle in the $[n]^d$ digraph has positive weight, and $\mathfrak{m}(x, y)$ is well-defined. Therefore, a shortest path from x to y is given by the rectilinear path obtained by decreasing the coordinates r with $x_r > y_r$ and increasing the coordinates r with $x_r < y_r$. A simple calculation yields

$$\mathfrak{m}(x, y) := \sum_{r:x_r>y_r} \sum_{t=y_r}^{x_r-1} u_r(t) - \sum_{r:x_r<y_r} \sum_{t=x_r}^{y_r-1} l_r(t) \tag{2}$$

If a function $f \in \mathscr{P}(\mathbf{B})$, then applying (1) on every edge of the path described above (the upper bound when we decrement a coordinate and the lower bound when we increment a coordinate), we get $f(x) - f(y) \le \mathfrak{m}(x, y)$ for any pair (x, y). Conversely, if $\forall x, y, f(x) - f(y) \le \mathfrak{m}(x, y)$, then considering neighboring pairs gives $f \in \mathscr{P}(\mathbf{B})$. This argument is encapsulated in the following lemma.

Lemma 2.2. $f \in \mathscr{P}(\mathbf{B})$ *iff* $\forall x, y \in [n]^d$, $f(x) - f(y) \le \mathfrak{m}(x, y)$.

When $\mathscr{P}(\mathbf{B})$ is monotonicity, $\mathfrak{m}(x, y) = 0$ if $x \prec y$ and ∞ otherwise. For the c-Lipschitz property, $\mathfrak{m}(x, y) = c\|x - y\|_1$. The salient properties of $\mathfrak{m}(x, y)$ are documented below and can be easily checked.

Lemma 2.3. $\mathfrak{m}(x, y)$ *satisfies the following properties.*

1. (Triangle Inequality.) *For any* x, y, z, $\mathfrak{m}(x, z) \le \mathfrak{m}(x, y) + \mathfrak{m}(y, z)$.

2. (Linearity.) *If* x, y, z *are such that for every* $1 \le r \le d$*, either* $x_r \le y_r \le z_r$ *or* $x_r \ge y_r \ge z_r$*, then* $\mathfrak{m}(x, z) = \mathfrak{m}(x, y) + \mathfrak{m}(y, z)$.

3. (Projection.) *Fix any dimension* r*. Let* x, y *be two points with* $x_r = y_r$*. Let* x' *and* y' *be the projection of* x, y *onto some other* r*-hyperplane. That is,* $x'_r = y'_r$*, and* $x'_j = x_j$*,* $y'_j = y_j$ *for* $j \ne r$*. Then,* $\mathfrak{m}(x, y) = \mathfrak{m}(x', y')$ *and* $\mathfrak{m}(x, x') = \mathfrak{m}(y, y')$*.*

Proof. $\mathfrak{m}(x, x) = 0$ follows since the RHS of (2) is empty. Triangle inequality holds because $\mathfrak{m}(x, y)$ is a shortest path weight. Linearity follows by noting $\sum_{t=y_r}^{x_r-1} u_r(t) = \sum_{t=y_r}^{z_r-1} u_r(t) + \sum_{t=z_r}^{x_r-1} u_r(t)$. For projection, note that if $x_r = y_r$, the RHS of (2) has no term corresponding to r. Thus, $\mathfrak{m}(x, y) = \mathfrak{m}(x', y')$. Suppose $x'_r > x_r$. Then, $\mathfrak{m}(x, x') = \sum_{t=x_r}^{x'_r} u_i(t) = \mathfrak{m}(y, y')$. A similar proof holds when $x'_r < x_r$. □

Henceforth, all we need is Lemma 2.2 and Lemma 2.3. We will interchangably use the terms $\mathscr{P}(\mathbf{B})$ and $\mathscr{P}(\mathfrak{m})$ where $\mathfrak{m}$ is as defined in (2). In fact, since $\mathbf{B}$ and therefore $\mathfrak{m}$ will be fixed in most of our discussion, we will simply use $\mathscr{P}$ including the parametrization wherever necessary.

Definition 2.4 (Violation Graph). *The violation graph of a function f with respect to property $\mathscr{P}$, denoted as $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P})$, has $[n]^d$ as vertices, and edge (x, y) if it forms a violation to $\mathscr{P}$, that is either $f(x) - f(y) > \mathfrak{m}(x, y)$ or $f(y) - f(x) > \mathfrak{m}(y, x)$.*

The triangle inequality of $\mathfrak{m}$ suffices to prove the following version of a classic lemma [FLN+02] relating the distance of a function to $\mathscr{P}$ to the vertex cover of the violation graph.

Lemma 2.5. *For any distribution $\mathscr{D}$ on $[n]^d$, any bounded-derivative property $\mathscr{P}$, and any function f, $\mathsf{dist}_{\mathscr{D}}(f, \mathscr{P}) = \min_X \mu_{\mathscr{D}}(X)$ where the minimum is over all vertex covers of $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P})$. Thus, if M is* any *maximal matching in $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P})$, then for the* uniform distribution, *$|M| \geq \mathsf{dist}_{\mathscr{U}}(f, \mathscr{P}) n^d/2$.*

3 The Dimension Reduction Theorem

For any combinatorial line ℓ in $[n]^d$, $f|_\ell : [n] \mapsto \mathbb{R}$ is f restricted to ℓ. It is natural to talk of $\mathscr{P}$ for any restriction of f, so $\mathsf{dist}_{\mathscr{D}_r}(f_{|\ell}, \mathscr{P})$ is well-defined for any r-line ℓ. For any $1 \leq r \leq d$, define the r-distance of the function:

$$\mathsf{dist}^r_{\mathscr{D}}(f, \mathscr{P}) := \mathbf{E}_{\ell \sim \mathscr{D}_{-r}}[\mathsf{dist}_{\mathscr{D}_r}(f_{|\ell}, \mathscr{P})] \tag{3}$$

Call a function f r-good if there are no violations along r-lines, that is, for any x and y on the same r-line, we have $f(x) - f(y) \leq \mathfrak{m}(x, y)$. Observe that $\mathsf{dist}^r_{\mathscr{D}}(f, \mathscr{P})$ is the minimum $\mu_{\mathscr{D}}$-mass of points on which f needs to be modified to make it r-good. The following is the optimal dimension reduction theorem which connects the r-distances to the real distance.

Theorem 3.1 (Dimension Reduction)**.** *For any function f, any bounded-derivative property $\mathscr{P}$, and any product distribution $\mathscr{D} = \prod_{1 \leq r \leq d} \mathscr{D}_i$,*

$$\sum_{r=1}^{d} \mathsf{dist}^r_{\mathscr{D}}(f, \mathscr{P}) \geq \mathsf{dist}_{\mathscr{D}}(f, \mathscr{P})/4.$$

(It can be easily shown that $\sum_{r=1}^{d} \mathsf{dist}^r_{\mathscr{D}}(f, \mathscr{P}) \leq \mathsf{dist}_{\mathscr{D}}(f, \mathscr{P})$, by simply putting the same 1D function of all, say, 1-lines.) We first prove the above theorem for the uniform distribution. Recall the violation graph $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P})$ whose edges are violation to $\mathscr{P}$. We define weights on the edges (x, y).

$$w(x, y) := \mathsf{max}(f(x) - f(y) - \mathfrak{m}(x, y), f(y) - f(x) - \mathfrak{m}(y, x)) \tag{4}$$

Note that $w(x, y) > 0$ for all edges in the violation graph. Let M be a maximum weight matching of minimum cardinality (MWmC). (Introduce an arbitrary tie-breaking rule to ensure this is unique.) A pair $(x, y) \in M$ is an *r-cross pair* if $x_r \neq y_r$. The following theorem (proof defered to §3.2) establishes the crucial structural result about these MWmC matchings in violated graphs of r-good functions.

Theorem 3.2 (No r-violations $\Rightarrow$ no r-cross pairs)**.** *Let f be an r-good function. Then there exists an MWmC matching M in $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P})$ with no r-cross pairs.*

We proceed with the proof of Theorem 3.1 over the uniform distribution starting with some definitions.

Definition 3.3 (*Hypergrid slices*). *Given an r-dimensional vector $\mathbf{a} \in [n]^r$, the a-*slice *is $S_{\mathbf{a}} := \{x \in [n]^d : x_j = \mathbf{a}_j,\ 1 \le j \le r\}$.*

Each $\mathbf{a}$-slice is a $(d-r)$-dimensional hypergrid, and the various $\mathbf{a}$-slices for $\mathbf{a} \in [n]^r$ partition $[n]^d$. Let $f_{|\mathbf{a}}$ denote the restriction of f to the slice $S_{\mathbf{a}}$. For two functions f, g we use $\Delta(f,g) := |\{x : f(x) \neq g(x)\}| = \mathsf{dist}_{\mathscr{U}}(f,g) \cdot n^d$. The following claim relates the sizes of MWmC matchings to $\Delta(f,g)$.

Claim 3.4. *Let $f, g : [n]^d \mapsto \mathbb{R}$. Let M and N be the MWmC matchings in the violation graphs for f and g, respectively. Then, $||M| - |N|| \le \Delta(f,g)$.*

Proof. The symmetric difference of M and N is a collection of alternating paths and cycles. $||M| - |N||$ is at most the number of alternating paths. Each alternating path must contain a point at which f and g differ, for otherwise we can improve either M or N, either in weight or cardinality. □

Define a sequence of $d+1$ matchings $(M_0, M_1, \ldots, M_d)$ in $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P})$ in non-increasing order of cardinality as follows. For $0 \le r \le d$, M_r is the MWmC matching in $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P})$ among matchings that *do not contain any i-cross pairs for $1 \le i \le r$*. By Lemma 2.5, we have $|M_0| \ge \mathsf{dist}_{\mathscr{U}}(f, \mathscr{P}) n^d/2$. The last matching M_d is empty and thus has cardinality 0.

Lemma 3.5. *For all $1 \le r \le d$, we have $|M_{r-1}| - |M_r| \le 2 \cdot \mathsf{dist}^r_{\mathscr{U}}(f, \mathscr{P}) \cdot n^d$.*

Adding the inequalities in the statement of Lemma 3.5 for all r, we get $\mathsf{dist}_{\mathscr{U}}(f, \mathscr{P}) n^d/2 \le |M_0| - |M_d| \le 2\sum_{r=1}^{d} \mathsf{dist}^r_{\mathscr{U}}(f, \mathscr{P}) \cdot n^d$. This completes the proof of Theorem 3.1 for the uniform distribution. Now we prove Lemma 3.5.

Proof. Since M_{r-1} has no j-cross pairs for $1 \le j \le r-1$, all pairs of M_{r-1} have both endpoints in the same slice $S_{\mathbf{a}}$ for some $\mathbf{a} \in [n]^{r-1}$. Thus, M_{r-1} partitions into sub-matchings in each $S_{\mathbf{a}}$. Let $M^{\mathbf{a}}_{r-1}$ be the pairs of M_{r-1} with both endpoints in slice $S_{\mathbf{a}}$, so $|M_{r-1}| = \sum_{\mathbf{a} \in [n]^{r-1}} |M^{\mathbf{a}}_{r-1}|$. Similarly, $M^{\mathbf{a}}_r$ is defined. Since M_r has no r-cross pairs either, $\forall \mathbf{a} \in [n]^{r-1}$, $|M^{\mathbf{a}}_r| = \sum_{i=1}^{n} |M_r^{(\mathbf{a} \circ i)}|$, where $(\mathbf{a} \circ i)$ is the r-dimensional vector obtained by concatenating i to the end of $\mathbf{a}$. Observe that for any $\mathbf{a} \in [n]^{r-1}$, $M^{\mathbf{a}}_{r-1}$ is an MWmC matching in $S_{\mathbf{a}}$ w.r.t. $f_{|\mathbf{a}}$. Furthermore, for any $i \in [n]$, $M_r^{(\mathbf{a} \circ i)}$ is an MWmC matching in $S_{(\mathbf{a} \circ i)}$ w.r.t. $f_{|(\mathbf{a} \circ i)}$. Let $f^{(r)}$ be the closest function to f with no violations along dimension r. By definition, $\Delta(f, f^{(r)}) = \mathsf{dist}^r(f, \mathscr{P}) \cdot n^d$. Now comes the crucial part of the proof. Fix $\mathbf{a} \in [n]^{r-1}$ and focus on the $\mathbf{a}$-slice $S_{\mathbf{a}}$. Since $f^{(r)}$ has no violations along the r-lines, neither does ${f^{(r)}}_{|\mathbf{a}}$. By Theorem 3.2, there exists an MWmC matching $N^{\mathbf{a}}$ in $S_{\mathbf{a}}$ w.r.t. ${f^{(r)}}_{|\mathbf{a}}$ which has no r-cross pairs. Therefore, $N^{\mathbf{a}}$ partitions as $N^{\mathbf{a}} = \bigcup_{i=1}^{n} N^{(\mathbf{a} \circ i)}$. Furthermore, each matching $N^{(\mathbf{a} \circ i)}$ is an MWmC matching in $S_{(\mathbf{a} \circ i)}$ with respect to the weights corresponding to the function $f^{(r)}_{|(\mathbf{a} \circ i)}$. Since $M^{\mathbf{a}}_{r-1}$ is an MWmC matching w.r.t. $f_{|\mathbf{a}}$ and $N^{\mathbf{a}}$ is an MWmC matching w.r.t. $f^{(r)}_{|\mathbf{a}}$ in $S_{\mathbf{a}}$, Claim 3.4 gives

$$|N^{\mathbf{a}}| \ge |M^{\mathbf{a}}_{r-1}| - \Delta(f_{|\mathbf{a}}, {f^{(r)}}_{|\mathbf{a}}) \tag{5}$$

Since $M_r^{(\mathbf{a} \circ i)}$ is an MWmC matching w.r.t. $f_{|(\mathbf{a} \circ i)}$ and $N^{(\mathbf{a} \circ j)}$ is an MWmC matching w.r.t. $f^{(r)}_{|(\mathbf{a} \circ i)}$ in $S_{(\mathbf{a} \circ i)}$, Claim 3.4 gives us $|M_r^{(\mathbf{a} \circ i)}| \ge |N^{(\mathbf{a} \circ i)}| - \Delta(f_{|(\mathbf{a} \circ i)}, {f^{(r)}}_{|(\mathbf{a} \circ i)})$. Summing over all $1 \le i \le n$,

$$|M^{\mathbf{a}}_r| \ge |N^{\mathbf{a}}| - \Delta(f_{|\mathbf{a}}, {f^{(r)}}_{|\mathbf{a}}) \tag{6}$$

Adding (5), (6) over all $\mathbf{a} \in [n]^{r-1}$, $|M_r| \geq |M_{r-1}| - 2\sum_{\mathbf{a}\in[n]^{r-1}} \Delta(f_{|\mathbf{a}}, {f^{(r)}}_{|\mathbf{a}}) = |M_{r-1}| - 2 \cdot \mathsf{dist}^r(f, \mathscr{P}) \cdot n^d$. □

3.1 Reducing from arbitrary product distributions

We reduce arbitrary product distributions to uniform distributions on what we call the bloated hypergrid. Assume without loss of generality that all $\mu_{\mathscr{D}_r}(j) = q_r(j)/N$, for some integers $q_r(j)$ and N. Consider the d-dimensional N-hypergrid $[N]^d$. There is a natural many-to-one mapping from $\Phi : [N]^d \mapsto [n]^d$ defined as follows. First fix a dimension r. Given an integer $1 \leq t \leq N$, let $\phi_r(t)$ denote the index $\ell \in [1, n]$ such that $\sum_{j<\ell} q_r(j) < t \leq \sum_{j\leq\ell} q_r(j)$. That is, partition $[N]$ into n contiguous segments of lengths $q_r(1), \ldots, q_r(n)$. Then $\phi_r(t)$ is the index of the segment where t lies. The mapping $\Phi : [N]^d \mapsto [n]^d$ is defined as

$$\Phi(x_1, x_2 \ldots, x_d) = (\phi_1(x_1), \phi_2(x_2), \ldots, \phi_{\mathfrak{m}}(x_d)) .$$

We use Φ^{-1} to define the set of preimages, so Φ^{-1} maps a point in $[n]^d$ to a 'cuboid' in $[N]^d$. Observe that for any $x \in [n]^d$,

$$|\Phi^{-1}(x)| = N^d \prod_{r=1}^{d} \mu_{\mathscr{D}_r}(x) = N^d \mu_{\mathscr{D}}(x). \tag{7}$$

Claim 3.6. *For any set $X \subseteq [n]^d$, define $Z \subseteq [N]^d$ as $Z := \bigcup_{x\in X} \Phi^{-1}(x)$. Then $\mu_{\mathscr{D}}(X) = \mu_{\mathscr{U}}(Z)$.*

Proof. The set $Z = \bigcup_{x\in X} \Phi^{-1}(x)$ is the union of all the preimages of Φ over the elements of X. Since preimages are disjoint, we get $|Z| = \sum_{x\in X} |\Phi^{-1}(x)| = N^d \mu_{\mathscr{D}}(X)$. Therefore, $\mu_{\mathscr{U}}(Z) = \mu_{\mathscr{D}}(X)$. □

Given $f : [n]^d \mapsto \mathbb{R}$, we define its extension $f_{\mathtt{ext}} : [N]^d \mapsto \mathbb{R}$:

$$f_{\mathtt{ext}}(x_1, \ldots, x_d) = f(\Phi(x_1, \ldots, x_d)). \tag{8}$$

Thus, $f_{\mathtt{ext}}$ is constant on the cuboids in the bloated hypergrid corresponding to a point in the original hypergrid. Define the following metric on $[N]^d$.

$$\text{For } x, y \in [N]^d, \quad \mathfrak{m}_{\mathtt{ext}}(x, y) = \mathfrak{m}(\Phi(x), \Phi(y)) \tag{9}$$

The following statements establish the utility of the bloated hypergrid, and the proof of the dimension reduction of f over $[n]^d$ w.r.t. $\mathscr{D}$ follows easily from these and the proof for the uniform distribution.

Lemma 3.7. *If $\mathfrak{m}$ satisfies the conditions of Lemma 2.3 over $[n]^d$, then so does $\mathfrak{m}_{\mathtt{ext}}$ over $[N]^d$.*

Proof. Consider $x, y, z \in [N]^d$. Triangle inequality and well-definedness immediately follow from the validity of $\mathfrak{m}$. Now for linearity. If $x_r \leq y_r \leq z_r$, then so is $\phi_r(x_r) \leq \phi_r(y_r) \leq \phi_r(z_r)$. Thus, $\Phi(x), \Phi(y), \Phi(z)$ satisfy linearity w.r.t. $\mathfrak{m}$. So, $\mathfrak{m}_{\mathtt{ext}}(x, z) = \mathfrak{m}(\Phi(x), \Phi(z)) = \mathfrak{m}(\Phi(x), \Phi(y)) + \mathfrak{m}(\Phi(y), \Phi(z)) = \mathfrak{m}_{\mathtt{ext}}(x, y) + \mathfrak{m}_{\mathtt{ext}}(y, z)$. Now for projection. Suppose $x_r = y_r$ and $x'_r = y'_r$. Note that $\Phi(x)$ and $\Phi(y)$ have same rth coordinate, and so do $\Phi(x')$ and $\Phi(y')$. Furthermore, $\Phi(x')$ (resp. $\Phi(y')$) is the projection of $\Phi(x)$ (resp. $\Phi(x)$). Thus we get $\mathfrak{m}_{\mathtt{ext}}(x, y) = \mathfrak{m}(\Phi(x), \Phi(y)) = \mathfrak{m}(\Phi(x'), \Phi(y')) = \mathfrak{m}_{\mathtt{ext}}(x', y')$, and similarly $\mathfrak{m}_{\mathtt{ext}}(x, x') = \mathfrak{m}_{\mathtt{ext}}(y, y')$. □

Theorem 3.8. $\mathsf{dist}_{\mathscr{D}}(f, \mathscr{P}(\mathfrak{m})) = \mathsf{dist}_{\mathscr{U}}(f_{\mathsf{ext}}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}}))$.

Proof. ($\geq$). Let $X \subseteq [n]^d$ be a vertex cover in $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P}(\mathfrak{m}))$ minimizing $\mu_{\mathscr{D}}(X)$. From Lemma 2.5, $\mathsf{dist}_{\mathscr{D}}(f, \mathscr{P}(\mathfrak{m})) = \mu_{\mathscr{D}}(X)$. We claim $Z = \bigcup_{x \in X} \Phi^{-1}(x)$ is a vertex cover of $\mathcal{G}_{\mathsf{viol}}(f_{\mathsf{ext}}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}}))$. This implies $\mathsf{dist}_{\mathscr{U}}(f_{\mathsf{ext}}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}})) \leq \mu_{\mathscr{U}}(Z) = \mu_{\mathscr{D}}(X) = \mathsf{dist}_{\mathscr{D}}(f, \mathscr{P}(\mathfrak{m}))$, where the first equality follows from Claim 3.6. Consider a violated pair (u, v) in this graph and so wlog $f_{\mathsf{ext}}(u) - f_{\mathsf{ext}}(v) > \mathfrak{m}_{\mathsf{ext}}(u, v)$. Hence, $f(\Phi(u)) - f(\Phi(v)) > \mathfrak{m}(\Phi(u), \Phi(v))$ implying $(\Phi(u), \Phi(v))$ is an edge in $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P}(\mathfrak{m}))$. Thus, either $\Phi(u)$ or $\Phi(v)$ lies in X implying either u or v lies in Z.($\leq$). Let $Z \subseteq [N]^d$ be a vertex cover in $\mathcal{G}_{\mathsf{viol}}(f_{\mathsf{ext}}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}}))$ minimizing $\mu_{\mathscr{U}}(Z)$. Therefore, $\mathsf{dist}_{\mathscr{U}}(f_{\mathsf{ext}}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}})) = \mu_{\mathscr{U}}(Z)$. Define $X \subseteq [n]^d$ as $X = \{x \in [n]^d : \Phi^{-1}(x) \subseteq Z\}$. Therefore, $Z \supseteq \bigcup_{x \in X} \Phi^{-1}(x)$ and from Claim 3.6 we get $\mu_{\mathscr{U}}(Z) \geq \mu_{\mathscr{D}}(X)$. It suffices to show that X is a vertex cover of $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P}(\mathfrak{m}))$. Consider a violated edge (x, y) in this graph such that $f(x) - f(y) > \mathfrak{m}(x, y)$. Suppose neither x nor y are in X. Hence, there exists $u \in \Phi^{-1}(x) \setminus Z$ and $v \in \Phi^{-1}(y) \setminus Z$. So $f_{\mathsf{ext}}(u) - f_{\mathsf{ext}}(v) = f(\Phi(u)) - f(\Phi(v)) = f(x) - f(y) > \mathfrak{m}(x, y) = \mathfrak{m}_{\mathsf{ext}}(\Phi(u), \Phi(v))$, implying (u, v) is a violation in $\mathcal{G}_{\mathsf{viol}}(f_{\mathsf{ext}}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}}))$. This contradicts the fact that Z is a vertex cover. □

Fix a dimension r and r-line ℓ. Abusing notation, let $\Phi^{-1}(\ell)$ denote the collection of r-lines in $[N]^d$ that are mapped to ℓ by Φ. Note that $|\Phi^{-1}(\ell)| = N^{d-1}\mu_{\mathscr{D}_{-r}}(\ell)$. A proof identical to one above yields the following theorem.

Theorem 3.9. *For any r-line,* $\mathsf{dist}_{\mathscr{D}_r}(f_{|\ell}, \mathscr{P}(\mathfrak{m})) = \mathsf{dist}_{\mathscr{U}_r}(f_{\mathsf{ext}|\ell'}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}}))$ *for all* $\ell' \in \Phi^{-1}(\ell)$.

Now we can complete the proof of Theorem 3.1.

$$\begin{aligned}
\mathsf{dist}^r_{\mathscr{D}}(f, \mathscr{P}(\mathfrak{m})) &= \sum_{r\text{-line } \ell} \mu_{\mathscr{D}_{-r}}(\ell) \cdot \mathsf{dist}_{\mathscr{D}_r}(f_{|\ell}, \mathscr{P}(\mathfrak{m})) \\
&= \frac{1}{N^{d-1}} \sum_{r\text{-line } \ell} |\Phi^{-1}(\ell)| \cdot \mathsf{dist}_{\mathscr{D}_r}(f_{|\ell}, \mathscr{P}(\mathfrak{m})) \\
&= \frac{1}{N^{d-1}} \sum_{r\text{-line } \ell} \sum_{\ell' \in \Phi^{-1}(\ell)} \mathsf{dist}_{\mathscr{U}_r}(f_{\mathsf{ext}|\ell'}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}})) \\
&= \mathbf{E}_{\ell' \sim \mathscr{U}_{-r}}[\mathsf{dist}_{\mathscr{U}_r}(f_{\mathsf{ext}|\ell'}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}}))] = \mathsf{dist}^r_{\mathscr{U}}(f_{\mathsf{ext}}, \mathscr{P}(\mathfrak{m}_{\mathsf{ext}})).
\end{aligned}$$

We can apply the dimension reduction to f_{ext} for property $\mathscr{P}(\mathfrak{m}_{\mathsf{ext}})$ over the uniform distribution. The proof of Theorem 3.1 for f follows directly.

3.2 No r-violations imply no r-cross pairs.

In this subsection we prove Theorem 3.2. This closely follows the techniques and proofs from [CS13a].

Theorem 3.2 (No r-violations $\Rightarrow$ no r-cross pairs)**.** *Let f be an r-good function. Then there exists an MWmC matching M in $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P})$ with no r-cross pairs.*

This requires the alternating path setup of [CS13a]. Recall the weight function $w(x, y) = \mathsf{max}(f(x) - f(y) - \mathfrak{m}(x, y), f(y) - f(x) - \mathfrak{m}(y, x))$ defined on pairs of the domain. Note that (x, y) is a violation iff $w(x, y) > 0$. Let M be a maximum weight minimum cardinality (MWmC) matching of $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P}(\mathfrak{m}))$ with the *minimum number* of r-cross pairs. Recall an r-cross pair

(x, y) has $x_r \neq y_r$. We will prove that this minimum value is 0. Let $\mathsf{cr}(M)$ be the set of r-cross pairs in M. Let $\mathsf{st}(M) := M \setminus \mathsf{cr}(M)$. For contradiction's sake, assume $\mathsf{cr}(M)$ is nonempty. Let $(x, y) \in \mathsf{cr}(M)$ be an arbitrary r-cross pair with $x_r = a$ and $y_r = b$ with $a \neq b$. Define matching $H := \{(u, v) : u_r = a, v_r = b, u_j = v_j, j \neq i\}$. This is a matching by projection between points with rth coordinate a and b. For convenience, we denote the points with rth coordinate a (resp. b) as the *a-plane* (resp. b-plane). Consider the alternating paths and cycles in $H \Delta\, \mathsf{st}(M)$. The vertex y is incident to only an H-pair, since $(x, y) \in \mathsf{cr}(M)$. Let $y = s_1, s_2, \ldots, s_t$ be the alternating path starting from y, collectively denoted by S. We let $s_0 := x$. The end of S, s_t, may be either M-unmatched or $\mathsf{cr}(M)$-matched. In the latter case, we define s_{t+1} to be such that $(s_t, s_{t+1}) \in \mathsf{cr}(M)$. For even i, (s_{i-1}, s_i) is an H-pair and (s_i, s_{i+1}) is an M-pair. We list out some basic claims about S.

Claim 3.10. *If strictly positive $j \equiv 0, 1 \mod 4$, then s_j is in the b-plane. Otherwise, s_j is in the a-plane.*

Claim 3.11. *For strictly positive even i, $f(s_{i-1}) - f(s_i) - \mathfrak{m}(s_{i-1}, s_i) \leq 0$ and $f(s_i) - f(s_{i-1}) - \mathfrak{m}(s_i, s_{i-1}) \leq 0$.*

Proof. Since f is r-good and H-pairs differ only in the rth coordinate, $w(s_{i-1}, s_i) \leq 0$ for all even i. The definition of $w(s_{i-1}, s_i)$ completes the proof. □

Claim 3.12. *For strictly positive $i \equiv 0 \mod 4$, $\mathfrak{m}(s_{i-1}, s_i) = \mathfrak{m}(s_2, s_1)$. For $i \equiv 2 \mod 4$, $\mathfrak{m}(s_i, s_{i-1}) = \mathfrak{m}(s_2, s_1)$.*

Proof. The point s_0 (which is x) lies in the a-plane. Hence, for any $i \equiv 2 \mod 4$, s_i lies in the b-plane. Similarly, for $i \equiv 0 \mod 4$, s_i lies in the a-plan. For strictly positive even i, (s_{i-1}, s_i) is an H-pair. An application of the projection property completes the proof. □

Claim 3.13. *For strictly positive even i, $\mathfrak{m}(s_i, s_{i+1}) = \mathfrak{m}(s_{i-1}, s_{i+2})$ and $\mathfrak{m}(s_{i+1}, s_i) = \mathfrak{m}(s_{i+2}, s_{i-1})$.*

Proof. Consider $\mathsf{st}(M)$-pair (s_i, s_{i+1}). Both points are on the same (a or b-)plane. Observe that s_{i-1} is the projection of s_i and s_{i+2} is the projection of s_{i+1} onto the other plane. Apply the projection property of d to complete the proof. □

Now we have all the ingredients to prove the theorem. The strategy is to find another matching M' such that either $w(M') > w(M)$ or $w(M') = w(M)$ and M' has strictly fewer cross pairs. Let us identify certain subsets of pairs to this end. For even k, define

$$E_-(k) := (s_0, s_1), (s_2, s_3), \ldots, (s_k, s_{k+1}) = \{(s_j, s_{j+1}) : j \text{ even}, 0 \leq j \leq k\}$$

These are precisely the $\mathsf{st}(M)$-pairs in S in the first k-steps. Note that $|E_-(k)| = k/2 + 1$. Now we define $E_+(k)$. In English: first pick pair (s_0, s_2); subsequently pick the first unpaired s_i and pair it with the next unpaired s_j of the *opposite* parity. More precisely, for even k,

$$E_+(k) := (s_0, s_2), (s_1, s_4), (s_3, s_6), \ldots, (s_{k-3}, s_k) = (s_0, s_2) \cup \{(s_{j-3}, s_j) : j \text{ even}, 4 \leq j \leq k\}$$

Note that $|E_+(k)| = k/2$. Wlog, assume that $w(x, y) = f(x) - f(y) - \mathfrak{m}(x, y)$. It turns out the weights of all other M-pairs in S are determined. We will assert that the pattern is as follows.

$$w(s_i, s_{i+1}) = \begin{cases} f(s_i) - f(s_{i+1}) - \mathfrak{m}(s_i, s_{i+1}) & \text{if } i \equiv 0 \mod 4 \\ f(s_{i+1}) - f(s_i) - \mathfrak{m}(s_{i+1}, s_i) & \text{if } i \equiv 2 \mod 4 \end{cases} \qquad (\clubsuit)$$

The following lemma determines the weights of all other M-edges in the alternating path S. Recall $(s_i, s_{i+1}) \in \mathsf{st}(M)$ for even i.

Lemma 3.14. *Suppose s_i exists. If (♣) holds for all even indices $< i$, then s_i is matched in M.*

Proof. Assume $i \equiv 2 \mod 4$. (The other case is analogous and omitted.) We prove by contradiction, so suppose s_i is not matched in M. We set $M' := M - E_-(i-2) + E_+(i)$. Note that M' is a valid matching, since s_i is not matched. We compare $w(M')$ and $w(M)$. By (♣), we can express $w(E_-(i-2))$ exactly.

$$\begin{aligned} w(E_-(i-2)) &= \sum_{j:\text{even},\ 0\le j\le i-2} w(s_j, s_{j+1}) \\ &= [f(s_0) - f(s_1) - \mathfrak{m}(s_0,s_1)] + [f(s_3) - f(s_2) - \mathfrak{m}(s_3,s_2)] + \\ &\quad [f(s_4) - f(s_5) - \mathfrak{m}(s_4,s_5)] + [f(s_7) - f(s_6) - \mathfrak{m}(s_7,s_6)] + \cdots \\ &\quad [f(s_{i-2}) - f(s_{i-1}) - \mathfrak{m}(s_{i-2},s_{i-1})] \end{aligned} \tag{10}$$

We lower bound $w(E_+(i))$. Since each individual weight term is a maximum of two expressions, we can choose either. We set the expression up to match $w(E_-(i-2))$ as best as possible.

$$\begin{aligned} w(E_+(i)) &\ge [f(s_0) - f(s_2) - \mathfrak{m}(s_0,s_2)] + [f(s_4) - f(s_1) - \mathfrak{m}(s_4,s_1)] + \\ &\quad [f(s_3) - f(s_6) - \mathfrak{m}(s_3,s_6)] + [f(s_8) - f(s_5) - \mathfrak{m}(s_8,s_5)] + \\ &\quad [f(s_{i-3}) - f(s_i) - \mathfrak{m}(s_{i-3},s_i)] \end{aligned} \tag{11}$$

Note that $w(M') - w(M) = w(E_+(i)) - w(E_-(i-2))$. Observe that any f term that occurs in both (10) and (11) has the same coefficient. By Claim 3.13, $\mathfrak{m}(s_3,s_2) = \mathfrak{m}(s_4,s_1)$, $\mathfrak{m}(s_4,s_5) = \mathfrak{m}(s_3,s_6)$, etc.

$$w(E_+(i)) - w(E_-(i-2)) \ge f(s_{i-1}) - f(s_i) - \mathfrak{m}(s_0,s_2) + \mathfrak{m}(s_0,s_1)$$

The points s_0 and s_1 lie is different planes, and $(s_1, s_2) \in H$. We can apply the linearity property to get $\mathfrak{m}(s_0,s_1) = \mathfrak{m}(s_0,s_2) + \mathfrak{m}(s_2,s_1)$. Plugging this in, applying Claim 3.11 and Claim 3.12 for i,

$$w(E_+(i)) - w(E_-(i-2)) \ge f(s_{i-1}) - f(s_i) + \mathfrak{m}(s_2,s_1) = -[f(s_i) - f(s_{i-1}) - \mathfrak{m}(s_i,s_{i-1})] \ge 0$$

Hence $w(M') \ge w(M)$. Note that $|M'| - |M| = |E_+(i)| - |E_-(i-2)| = i/2 - ((i-2)/2 + 1) = 0$. Finally, observe that $E_+(i)$ has no r-cross pairs, but $E_-(i-2)$ has one (pair (s_0, s_1)). This contradicts the choice of M as a MWmC matching with the least r-cross pairs. □

Claim 3.15. *If (♣) holds for all even indices $< i$, then $s_0, s_1, \ldots, s_{i+1}$ are all distinct.*

Proof. (This is trivial if $i < t$. The non-trivial case if when S ends as s_i.) The points $s_1, \ldots, s_i$ are all distinct. If $s_i \neq x$, the claim holds. So assume $s_i = x = s_0$. By Claim 3.10, $i \equiv 2 \mod 4$. Replace pairs $A = \{(s_0,s_1), (s_{i-2}, s_{i-1})\}$ by (s_{i-2}, s_1). Note that $\mathfrak{m}(s_0,s_1) = \mathfrak{m}(s_0,s_{i-1}) + \mathfrak{m}(s_{i-1},s_1)$. By (♣),

$$\begin{aligned} w(A) &= [f(s_0) - f(s_1) - \mathfrak{m}(s_0,s_1)] + [f(s_{i-2}) - f(s_{i-1}) - \mathfrak{m}(s_{i-2},s_{i-1})] \\ &= [f(s_{i-2}) - f(s_1) - \mathfrak{m}(s_{i-2},s_{i-1}) - \mathfrak{m}(s_{i-1},s_1)] + [f(s_0) - f(s_{i-1}) - \mathfrak{m}(s_0,s_{i-1})] \\ &\le [f(s_{i-2}) - f(s_1) - \mathfrak{m}(s_{i-2},s_1)] \le w(s_{i-2},s_1) \end{aligned}$$

The total number of pairs has decreased, so we complete the contradiction. □

Lemma 3.16. *Suppose s_i exists. If (♣) holds for all even indices $< i$, then (♣) holds for i.*

Proof. We prove by contradiction, so (♣) is false for i. (Again, assume $i \equiv 2 \mod 4$. The other case is omitted.) By Claim 3.15, $E_+(i-2) \cup (s_{i-3}, s_{i+1})$ is a valid set of matched pairs. Let $M' := M - E_-(i) + (E_+(i-2) \cup (s_{i-3}, s_{i+1}))$. Observe that $|M'| = |M| - 1$ and the vertices s_{i-1} and s_i are left unmatched in M'. By (♣) for even indices $< i$ and the opposite of (♣) for i,

$$\begin{aligned} w(E_-(i)) \quad = \quad & [f(s_0) - f(s_1) - \mathfrak{m}(s_0, s_1)] + [f(s_3) - f(s_2) - \mathfrak{m}(s_3, s_2)] + \\ & [f(s_4) - f(s_5) - \mathfrak{m}(s_4, s_5)] + [f(s_7) - f(s_6) - \mathfrak{m}(s_7, s_6)] + \cdots \\ & [f(s_{i-2}) - f(s_{i-1}) - \mathfrak{m}(s_{i-2}, s_{i-1})] + [f(s_i) - f(s_{i+1}) - \mathfrak{m}(s_i, s_{i+1})] \end{aligned} \tag{12}$$

We stress that the last weight is "switched". We lower bound $w(E_+(i-2) \cup (s_{i-3}, s_{i+1}))$ appropriately.

$$\begin{aligned} w(E_+(i-2) \cup (s_{i-3}, s_{i+1})) \quad \geq \quad & [f(s_0) - f(s_2) - \mathfrak{m}(s_0, s_2)] + [f(s_4) - f(s_1) - \mathfrak{m}(s_4, s_1)] + \\ & [f(s_3) - f(s_6) - \mathfrak{m}(s_3, s_6)] + [f(s_8) - f(s_5) - \mathfrak{m}(s_8, s_5)] + \cdots \\ & [f(s_{i-7}) - f(s_{i-4}) - \mathfrak{m}(s_{i-7}, s_{i-4})] + [f(s_{i-2}) - f(s_{i-5}) - \mathfrak{m}(s_{i-2}, s_{i-5})] + \\ & [f(s_{i-3}) - f(s_{i+1}) - \mathfrak{m}(s_{i-3}, s_{i+1})] \end{aligned} \tag{13}$$

As before, we subtract (12) from (13). All function terms from (13) cancel out. By Claim 3.13, all $\mathfrak{m}$-terms except the first and last cancel out.

$$w(M') - w(M) \geq f(s_{i-1}) - f(s_i) - \mathfrak{m}(s_0, s_2) - \mathfrak{m}(s_{i-3}, s_{i+1}) + \mathfrak{m}(s_0, s_1) + \mathfrak{m}(s_{i-2}, s_{i-1}) + \mathfrak{m}(s_i, s_{i+1})$$

By linearity, $\mathfrak{m}(s_0, s_1) = \mathfrak{m}(s_0, s_2) + \mathfrak{m}(s_2, s_1)$. Furthermore, by Claim 3.12, $\mathfrak{m}(s_2, s_1) = \mathfrak{m}(s_i, s_{i-1})$. By Claim 3.13, $\mathfrak{m}(s_{i-2}, s_{i-1}) = \mathfrak{m}(s_{i-3}, s_i)$. By triangle inequality, $-\mathfrak{m}(s_{i-3}, s_{i+1}) + \mathfrak{m}(s_{i-3}, s_i) + \mathfrak{m}(s_i, s_{i+1}) \geq 0$. Putting it all together and applying Claim 3.11,

$$w(M') - w(M) \geq -[f(s_i) - f(s_{i-1}) - \mathfrak{m}(s_i, s_{i-1})] \geq 0$$

So M' has at least the same weight but lower cardinality than M. Contradiction. □

Lemma 3.17. *Suppose s_i exists. If (♣) holds for all even indices $< i$, then s_i is matched in $\mathsf{st}(M)$.*

Proof. Suppose not. (Again, assume $i \equiv 2 \mod 4$.) By Lemma 3.14, s_i is matched in M, so $(s_i, s_{i+1}) \in \mathsf{cr}(M)$. We set $M' = M - E_-(i) + (E_+(i) \cup (s_{i-1}, s_{i+1}))$. By Claim 3.15, M' is a valid matching. We have $|M'| = |M|$. M has two r-cross pairs (s_0, s_1) and (s_i, s_{i+1}), but M' has at most one (s_{i-1}, s_{i+1}). It suffices to show that $w(M') \geq w(M)$ to complete the contradiction. By Lemma 3.16 and (♣),

$$\begin{aligned} w(E_-(i)) \quad = \quad & [f(s_0) - f(s_1) - \mathfrak{m}(s_0, s_1)] + [f(s_3) - f(s_2) - \mathfrak{m}(s_3, s_2)] + \\ & [f(s_4) - f(s_5) - \mathfrak{m}(s_4, s_5)] + [f(s_7) - f(s_6) - \mathfrak{m}(s_7, s_6)] + \cdots \\ & [f(s_{i-2}) - f(s_{i-1}) - \mathfrak{m}(s_{i-2}, s_{i-1})] + [f(s_{i+1}) - f(s_i) - \mathfrak{m}(s_{i+1}, s_i)] \end{aligned}$$

$$\begin{aligned} w(E_+(i) \cup (s_{i-1}, s_{i+1})) \quad \geq \quad & [f(s_0) - f(s_2) - \mathfrak{m}(s_0, s_2)] + [f(s_4) - f(s_1) - \mathfrak{m}(s_4, s_1)] + \\ & [f(s_3) - f(s_6) - \mathfrak{m}(s_3, s_6)] + [f(s_8) - f(s_5) - \mathfrak{m}(s_8, s_5)] + \cdots \\ & [f(s_{i-3}) - f(s_i) - \mathfrak{m}(s_{i-3}, s_i)] + [f(s_{i+1}) - f(s_{i-1}) - \mathfrak{m}(s_{i+1}, s_{i-1})] \end{aligned}$$

All function terms and all but the first and last $\mathfrak{m}$-terms cancel out. The second inequality below holds by linearity and triangle inequality. The last equality is an application of Claim 3.12.

$$\begin{aligned} w(M') - w(M) &\geq \mathfrak{m}(s_0, s_1) - \mathfrak{m}(s_0, s_2) + \mathfrak{m}(s_{i+1}, s_i) - \mathfrak{m}(s_{i+1}, s_{i-1}) \\ &\geq \mathfrak{m}(s_2, s_1) - \mathfrak{m}(s_i, s_{i-1}) = 0 \end{aligned}$$

□

Finally, we prove Theorem 3.2.

Proof. We started with a MWmC matching M with the minimum number of r-cross pairs. If there exists at least one such cross pair (x, y), we can define the alternating path sequence S. Wlog, we assumed (♣) holds for $i = 0$. Applications of Lemma 3.16 and Lemma 3.17 imply that S can never terminate. Contradiction. □

4 Search Trees and Bounded Derivative Property Testing.

As a result of dimension reduction, we can focus on designing testers for the line $[n]$. Our analysis is simple, but highlights the connection between bounded-derivative property testing and optimal search trees.

4.1 Testers for the Line $[n]$.

Let T be *any* binary search tree (BST) with respect to the totally ordered domain $[n]$. Every node of T is labeled with a unique entry in $[n]$, and the left (resp. right) child, if it exists, has a smaller (resp. larger) entry. The *depth* of a node v in the tree T, denoted as $\mathsf{depth}_T(v)$, is the number of *edges* on its path to the root. So the root has depth 0. Given a distribution $\mathscr{D}$ on $[n]$, the expected depth of T w.r.t. $\mathscr{D}$ is denoted as $\Delta(T; \mathscr{D}) = \mathbf{E}_{v \sim \mathscr{D}}[\mathsf{depth}_T(v)]$. The depth of the optimal BST w.r.t. $\mathscr{D}$ is denoted by $\Delta^*(\mathscr{D})$. It has long been observed that the transitivity of violations is the key property required for monotonicity testing on $[n]$ [BRW05, EKK+00, ACCL07, JR11]. We distill this argument down to a key insight: Given *any* BST T, there exists the following tester BST(T) for $\mathscr{P}$ on the line.

BST Tester (T)
1. Sample $v \sim \mathscr{D}$.
2. If v is the root of T, do nothing.
3. Else, query $f(u)$ for all vertices lying on the path from v to root (including the root and v).
4. Reject if any pair of these vertices form a violation to $\mathscr{P}$.

It is clear that the tester never rejects a function satisfying $\mathscr{P}$. (To connect with previous work, observe that the list of ancestor-descendant pairs forms a 2-Transitive Closure spanner [BGJ+09].)

Lemma 4.1. *For any bounded derivative property $\mathscr{P}$,* $\Pr[\textit{BST tester rejects}] \geq \mathsf{dist}_{\mathscr{D}}(f, \mathscr{P})$.

Proof. Let X be the set of non-root nodes v of T with the following property: (u, v) is a violation to $\mathscr{P}$ for some node u on the path from v to the root of T. The probability of rejection of the BST tester is precisely $\mu_{\mathscr{D}}(X)$. We claim that X is a vertex cover of $\mathcal{G}_{\mathsf{viol}}(f, \mathscr{P})$ which proves the lemma using Lemma 2.5. Pick any violation (x, y) and assume without loss of generality $f(x) - f(y) > \mathfrak{m}(x, y)$. Let z be the lowest common ancestor of x and y in T. By the BST property, either $x < z < y$ of $x > z > y$. By the linearity property of $\mathfrak{m}$, we get $\mathfrak{m}(x, y) = \mathfrak{m}(x, z) + \mathfrak{m}(z, y)$. This implies either $f(x) - f(z) > \mathfrak{m}(x, z)$ or $f(z) - f(y) > \mathfrak{m}(z, y)$, that is, either (x, z) or (y, z) is a violation implying one of them is in X. □

Lemma 4.2. *For any BST T, there is a $24\varepsilon^{-1}\Delta(T; \mathscr{D})$-query line monotonicity-tester.*

Proof. The expected number of queries made by the BST tester is $\sum_{v:\text{ non-root}} \Pr[v] \cdot (\mathsf{depth}_T(v) + 1) = (1 - \Pr[\text{root}]) + \Delta(T; \mathscr{D}) \leq 2 \cdot \Delta(T; \mathscr{D})$.

$$\sum_{v:\text{ non-root}} \Pr[v] \cdot (\mathsf{depth}_T(v) + 1) = (1 - \Pr[\text{root}]) + \Delta(T; \mathscr{D}) \leq 2 \cdot \Delta(T; \mathscr{D})$$

The expected depth is at least $(1 - \Pr[\text{root}])$ since non-roots have depth at least 1. To get a bonafide tester with deterministic query bounds, run the BST tester $2/\varepsilon$ times, aborting (and accepting) if the total number of queries exceeds $24\Delta(T; \mathscr{D})/\varepsilon$. The expected total number of queries is at most $4\Delta(T; \mathscr{D})/\varepsilon$. By Markov's inequality, the probability that the tester aborts is $\leq 1/6$. By Lemma 4.1, if $\mathsf{dist}_{\mathscr{D}}(f; \mathscr{P}) > \varepsilon$, the probability that this tester does not find a violation is at most $(1 - \varepsilon)^{2/\varepsilon} \leq 1/6$. With probability $\geq (1 - 1/6 - 1/6) = 2/3$, the tester rejects an ε-far function. □

Choose T to be the optimal BST to get the following theorem.

Theorem 4.3. *There exists a $24\varepsilon^{-1}\Delta^*(\mathscr{D})$-query tester for any bounded derivative property over the line.*

Note that once the tree is fixed, the BST tester only needs random samples from the distribution. Pick T to be the balanced binary tree of depth $O(\log n)$ to get a *distribution-free* tester.

Theorem 4.4. *There exists a $24\varepsilon^{-1} \log n$-query* distribution free *tester for any bounded-derivative property over the line.*

4.2 Testers for the Hypergrid.

Given a series of BSTs $T_1, T_2, \ldots, T_d$ corresponding to each dimension, we have the following hypergrid BST tester.

Hypergrid BST Tester $(T_1, T_2, \ldots, T_d)$
1. Sample $x \sim \mathscr{D}$.
2. Choose dimension r u.a.r. and let ℓ be the r-line through x.
3. Run **BST Tester**(T_r) on $f_{|\ell}$.

Lemma 4.5. *For any set of BSTs $T_1, T_2, \ldots, T_d$, the probability of rejection is at least $\mathsf{dist}_{\mathscr{D}}(f, \mathscr{P})/4d$.*

Proof. Condition on an r-line being chosen. The probability distribution over r-lines for this tester is $\mathscr{D}_{-r}$. By Lemma 4.1, the rejection probability is at least $\mathbf{E}_{\ell \sim \mathscr{D}_{-r}}[\mathsf{dist}_{\mathscr{D}_r}(f_{|\ell}, \mathscr{P})] = \mathsf{dist}^r_{\mathscr{D}}(f, \mathscr{P})$. The overall rejection probability is at least $\sum_{r=1}^{d} \frac{\mathsf{dist}^i_{\mathscr{D}}(f,\mathscr{P})}{d} \geq \frac{\mathsf{dist}_{\mathscr{D}}(f,\mathscr{P})}{4d}$, by Theorem 3.1. ☐

The expected number of queries made by this procedure is at most $\frac{1}{d} \cdot \sum_{r=1}^{d} 2\Delta(T_r; \mathscr{D}_r)$. Repeating it $O(d/\varepsilon)$ times to get the desired tester. The proof of the following is identical to that of Lemma 4.2 and is omitted.

Lemma 4.6. *For any collections of BSTs $(T_1, \ldots, T_d)$, there is a $100\varepsilon^{-1} \sum_{i=1}^{d} \Delta(T; \mathscr{D})$-query tester for any bounded derivative property.*

As in the case of the line we get the following as corollaries.

Theorem 1.3. [Main upper bound] *Consider functions $f : [n]^d \mapsto \mathbb{R}$. Let $\mathbf{B}$ be a bounding family and $\mathscr{D}$ be a product distribution. There is a tester for $\mathscr{P}(\mathbf{B})$ w.r.t. $\mathscr{D}$ making $100\varepsilon^{-1}\Delta^*(\mathscr{D})$ queries.*

Theorem 1.7. *Consider functions $f : [n]^d \mapsto \mathbb{R}$. There is a distribution-free (non-adaptive, one-sided) tester for $\mathscr{P}(\mathbf{B})$ w.r.t. $\mathscr{D}$ making $100\varepsilon^{-1} d \log n$ queries.*

The upper bound $\sum_{r=1}^{d} \Delta^*(\mathscr{D}_r)$ is at most $H(\mathscr{D})$, but can be much smaller, and it is clearest in the case of the hypercube. In the hypercube, each $\mathscr{D}_r$ is given by $(\mu_r, 1-\mu_r)$. Set $\theta_r := \min(\mu_r, 1-\mu_r)$. The optimal BST places the point of larger mass on the root and has expected depth θ_r.

Corollary 1.5. *Consider functions $f : \{0,1\}^d \mapsto \mathbb{R}$. Monotonicity testing over any product distribution $\mathscr{D} = \prod_{r=1}^{d} \mathscr{D}_r$, where each $\mathscr{D}_r = (\mu_r, 1-\mu_r)$, can be done with $100\varepsilon^{-1} \sum_{r=1}^{d} \min(\mu_r, 1-\mu_r)$ queries.*

It is instructive to open up this tester. It samples a point x from the distribution $\mathscr{D}$ and picks a dimension r uniformly at random. With probability θ_r, it queries both endpoints of $(x, x \oplus \mathbf{e}_r)$. With probability $(1-\theta_r)$, it does *nothing*. This process is repeated $O(d/\varepsilon)$ times. When $\theta_r = \mu_r = 1/2$, this is the standard edge tester.

5 Lower Bounds

We prove that the upper bounds of §4 are tight up to the dependence on the distance parameter ε. As alluded to in §1.3, we can only prove lower bounds for *stable* product distributions. These are distributions where small perturbations to the mass function do not change Δ^* drastically.

Definition 5.1 (Stable Distributions). *A product distribution $\mathscr{D}$ is said to be (ε, ρ)-stable if for all product distributions $\mathscr{D}'$ with $||\mathscr{D} - \mathscr{D}'||_{\mathsf{TV}} \leq \varepsilon$, $\Delta^*(\mathscr{D}') \geq \rho\Delta^*(\mathscr{D})$.*

The uniform distribution on $[n]^d$ is $(\varepsilon, 1-o(1))$-stable, for any constant $\varepsilon < 1$. The Gaussian distribution also shares the same stability. An example of an unstable distribution is the following. Consider $\mathscr{D}$ on $[n]$, where the probability on the first $k = \log n$ elements is $(1-\varepsilon)/k$, and is $\varepsilon/(n-k)$ for all other elements. Let $\mathscr{D}'$ have all its mass uniformly spread on the first k elements. We have $||\mathscr{D} - \mathscr{D}'||_{\mathsf{TV}} = \varepsilon$ but $\Delta^*(\mathscr{D}) \approx \varepsilon \log n$ and $\Delta^*(\mathscr{D}') \approx \log k = \log\log n$.

Theorem 1.6. [Main lower bound] *For any parameter ε, there exists $\varepsilon' = \Theta(\varepsilon)$ such that for any bounding family $\mathbf{B}$ and (ε', ρ)-stable, product distribution $\mathscr{D}$, any (even adaptive, two-sided) tester for $\mathscr{P}(\mathbf{B})$ w.r.t. $\mathscr{D}$ with proximity parameter ε requires $\Omega(\rho\Delta^*(\mathscr{D}))$ queries.*

5.1 Reduction from monotonicity to bounded-derivative property

Consider a function $f : [n]^d \mapsto [R]$ with where $R \in \mathbb{N}$. Let $\mathfrak{m}$ be the distance function obtained by bounding family $\mathbf{B}$. We let $\mathbf{0} \in [n]^d$ be $(0, 0, \ldots, 0)$. We use $\prec$ to denote the natural partial order in $[n]^d$, and let $\texttt{hcd}(x, y)$ be the highest common descendant of $x, y \in [n]^d$. We first prove an observation about triangle equality.

Observation 5.2. *If* $\mathfrak{m}(\mathbf{0}, x) + \mathfrak{m}(x, y) = \mathfrak{m}(\mathbf{0}, y)$*, then* $x \prec y$*.*

Proof. By linearity, $\mathfrak{m}(x, y) = \mathfrak{m}(x, \texttt{hcd}(x, y)) + \mathfrak{m}(\texttt{hcd}(x, y), y)$. Since $\texttt{hcd}(x, y) \prec x$, by linearity again, $\mathfrak{m}(\mathbf{0}, x) = \mathfrak{m}(\mathbf{0}, \texttt{hcd}(x, y)) + \mathfrak{m}(\texttt{hcd}(x, y), x)$. (Similarly for y.) Putting it all into the 'if' condition,

$$\mathfrak{m}(\mathbf{0}, \texttt{hcd}(x, y)) + \mathfrak{m}(\texttt{hcd}(x, y), x) + \mathfrak{m}(x, \texttt{hcd}(x, y)) + \mathfrak{m}(\texttt{hcd}(x, y), y) = \mathfrak{m}(\mathbf{0}, \texttt{hcd}(x, y)) + \mathfrak{m}(\texttt{hcd}(x, y), y)$$

This yields $\mathfrak{m}(\texttt{hcd}(x, y), x) + \mathfrak{m}(x, \texttt{hcd}(x, y)) = 0$. Suppose $\texttt{hcd}(x, y) \neq x$. The length (in terms of $\mathbf{B}$) of the path from $\texttt{hcd}(x, y)$ to x involves a sum of $u_i(t)$ terms, and the reverse path involves corresponding $-l_i(t)$ terms. Since $u_i(t) > l_i(t)$, the total path length from $\texttt{hcd}(x, y)$ to x and back is strictly positive. Therefore, $\texttt{hcd}(x, y) = x$ and $x \prec y$. □

Let U be the set of incomparable (ordered) pairs in $[n]^d$. Define $\delta := \min_{(x,y)\in U}\{\mathfrak{m}(\mathbf{0}, x) + \mathfrak{m}(x, y) - \mathfrak{m}(\mathbf{0}, y)\}$. By Observation 5.2, $\delta > 0$. Define

$$g(x) := \frac{\delta}{2R} \cdot f(x) - \mathfrak{m}(\mathbf{0}, x)$$

Lemma 5.3. $\mathsf{dist}_{\mathscr{D}}(g, \mathscr{P}) = \mathsf{dist}_{\mathscr{D}}(f, \texttt{MON})$.

Proof. We show that (u, v) violates $\mathscr{P}(\mathfrak{m})$ of g iff it violates monotonicity of f. First, the 'only if' case. Assume $g(u) - g(v) > \mathfrak{m}(u, v)$. Plugging in the expression for $g(\cdot)$ and rearranging,

$$\frac{\delta}{2R}(f(u) - f(v)) > \mathfrak{m}(\mathbf{0}, u) + \mathfrak{m}(u, v) - \mathfrak{m}(\mathbf{0}, v)$$

By triangle inequality on the RHS, $f(u) > f(v)$. Note that $f(u) - f(v) \leq R$ so $\frac{\delta}{2R}(f(u) - f(v)) \leq \delta/2$. So $\delta/2 > \mathfrak{m}(\mathbf{0}, u) + \mathfrak{m}(u, v) - \mathfrak{m}(\mathbf{0}, v)$. By choice of δ, the RHS must be zero. By Observation 5.2, $u \prec v$, and (u, v) is a violation to monotonicity of f. Now the 'only if' case, so $u \prec v$ and $f(u) > f(v)$. Note that $\mathfrak{m}(\mathbf{0}, v) = \mathfrak{m}(\mathbf{0}, u) + \mathfrak{m}(u, v)$. We deduce that (u, v) is also a violation to $\mathscr{P}(\mathfrak{m})$ for g.

$$g(u) - g(v) = \frac{\delta}{2R}(f(u) - f(v)) + \mathfrak{m}(\mathbf{0}, v) - \mathfrak{m}(\mathbf{0}, u) = \frac{\delta}{2R}(f(u) - f(v)) + \mathfrak{m}(u, v) > \mathfrak{m}(u, v)$$

□

Our main reduction theorem is the following.

Theorem 5.4. *Fix domain* $[n]^d$ *and a product distribution* $\mathscr{D}$*. Suppose there exists a* Q*-query tester for testing a bounded-derivative property* $\mathscr{P}$ *with distance parameter* ε*. Then there exists a* $Q + 10/\varepsilon$*-query tester for monotonicity for functions* $f : [n]^d \mapsto \mathbb{N}$ *over* $\mathscr{D}$ *with distance parameter* 2ε*.*

Proof. The monotonicity tester first queries $10/\varepsilon$ points of $[n]^d$, each i.i.d. from $\mathscr{D}$. Let the maximum f-value among these be this M. Consider the truncated function $f' : [n]^d \mapsto [M]$, where $f'(x) = M$ if $f(x) \geq M$ and $f'(x) = f(x)$ otherwise. If f is monotone, f' is monotone. Note that $\mathsf{dist}_{\mathscr{D}}(f, f') < \varepsilon$. So if f is 2ε-far from monotone, f' is ε-far from monotone. We can apply the $\mathscr{P}(\mathbf{B})$ tester on the function g obtained from Lemma 5.3. □

5.2 Monotonicity Lower Bound Framework.

The lower bound for monotonicity testing goes by the proof strategy set up in [CS13b]. This is based on arguments in [Fis04, CS13b] that reduce general testers to comparison-based testers. We encapsulate the main approach in the following theorem, proven implicitly in [CS13b]. (We use $\texttt{MON}$ to denote the monotonicity property.)

Theorem 5.5. *Fix domain $[n]^d$, distribution $\mathscr{D}$, proximity parameter ε, and positive integer L possibly depending on $\mathscr{D}$ and ε. A pair (x, y) distinguishes function g from h if $h(x) < h(y)$ and $g(x) > g(y)$. Suppose there is a collection of 'hard' functions $h, g_1, \ldots, g_L : [n]^d \mapsto \mathbb{N}$ such that*

- *The function h is monotone.*
- *Every $\mathsf{dist}_{\mathscr{D}}(g_i, \texttt{MON}) \geq \varepsilon$.*
- *Pairs in any set $Q \subset [n]^d$, can distinguish at most $|Q|$ of the g_i's from h.*

Then any (even adaptive, two-sided) monotonicity tester w.r.t. $\mathscr{D}$ for functions $f : [n]^d \mapsto \mathbb{N}$ with distance parameter ε must make $\Omega(L)$ queries.

In §5.3 and §5.4, we first describe hard functions for the line and the hypercube domain, respectively. The general hypergrid is addressed in §5.5.

5.3 The Line

Theorem 5.6. *Fix a parameter ε. If $\mathscr{D}$ is $(2\varepsilon, \rho)$-stable, then any ε-monotonicity tester w.r.t. $\mathscr{D}$ for functions $f : [n]^d \mapsto \mathbb{N}$ requires $\Omega(\rho\Delta^*(\mathscr{D}))$ queries.*

Not surprisingly, the lower bound construction is also based on BSTs. We specifically use the *median BST* [Meh75]. When $n = 1$, then the tree is the singleton. For a general n, let $t \in [n]$ be the smallest index such that $\mu(\{1, \cdots, t\}) \geq 1/2$ (henceforth, in this section, we use μ to denote $\mu_{\mathscr{D}}$). The root of T is t. Recur the construction on the intervals $[1, t-1]$ and $[t+1, n]$. By construction, the probability mass of any subtree together with its parent is greater than the probability mass of the sibling subtree. This *median property* will be utilized later. We follow the framework of Theorem 5.5 to construct a collection of hard functions. The monotone function h can be anything; $h(i) = 3i$ works. We will construct a function g_j ($j \geq 1$) for each non-root level of the median BST. Consider the nodes at depth $j-1$ (observe the use of $j-1$, and not j). Each of these corresponds to an interval, and we denote this sequence of intervals by $\mathsf{I}_j^1, \mathsf{I}_j^2, \ldots$. (Because internal nodes of the tree are also elements in $[n]$, there are gaps between these intervals.) Let $L_{\geq j} := \{x : \mathsf{depth}_T(x) \geq j\}$ be the nodes at depth j and higher. We have the following simple claim.

Claim 5.7. *I_j^k can be further partitioned into $\mathsf{I}_j^{k,\mathsf{left}}$ and $\mathsf{I}_j^{k,\mathsf{right}}$ such that $\sum_k \min\left(\mu(\mathsf{I}_j^{k,\mathsf{left}}), \mu(\mathsf{I}_j^{k,\mathsf{right}})\right) \geq \frac{\mu(L_{\geq j})}{2}$.*

Proof. Consider the node u_k corresponding I_j^k, and let the nodes in the left and right subtrees be S_ℓ and S_r. If $\mu(S_\ell) \le \mu(S_r)$, then $\mathsf{I}_j^{k,\mathsf{left}} = S_\ell \cup u_k$ and $\mathsf{I}_j^{k,\mathsf{right}} = S_r$. Otherwise, $\mathsf{I}_j^{k,\mathsf{left}} = S_\ell$ and $\mathsf{I}_j^{k,\mathsf{right}} = u_k \cup S_r$. By the median property of the BST, $\min(\mu(\mathsf{I}_j^{k,\mathsf{left}}), \mu(\mathsf{I}_j^{k,\mathsf{right}})) = \mathsf{max}(\mu(S_\ell), \mu(S_r)) \ge (\mu(S_\ell) + \mu(S_r))/2$. ☐

We describe the non-monotone g_j's and follow up with some claims. Let $\mathtt{lca}(x, y)$ denote the least common ancestor of x and y in T.

$$g_j(x) = \begin{cases} 2x & \text{if } x \notin \bigcup_k \mathsf{I}_j^k \\ 2x + 2(b-m) + 1 & \text{if } x \in \mathsf{I}_j^{k,\mathsf{left}} = [a, m], \quad \text{where } \mathsf{I}_j^k = [a, b]. \\ 2x - 2(m-a) - 1 & \text{if } x \in \mathsf{I}_j^{k,\mathsf{right}} = [m+1, b], \text{ where } \mathsf{I}_j^k = [a, b]. \end{cases} \tag{14}$$

Claim 5.8. (i) $\mathsf{dist}_{\mathscr{D}}(g_j, \mathtt{MON}) \ge \frac{\mu(L_{\ge j})}{2}$. (ii) *If (x, y) distinguishes g_j from h, then $\mathtt{lca}(x, y)$ lies in level $(j-1)$.*

Proof. All elements in $\mathsf{I}_j^{k,\mathsf{left}}$ are in violation with all elements in $\mathsf{I}_j^{k,\mathsf{right}}$ for all k. To see this, let $x \in \mathsf{I}_j^{k,\mathsf{left}}$ and $y \in \mathsf{I}_j^{k,\mathsf{right}}$, and so $x \prec y$. Denote $\mathsf{I}_j^k = [a, b]$,

$$g_j(x) - g_j(y) = 2x + 2(b-m) + 1 - 2y + 2(m-a) + 1 = 2(x-a) + 2(b-y) + 2 > 0$$

The vertex cover of the violation graph of g_i has mass at least $\sum_k \min(\mu(\mathsf{I}_j^{k,\mathsf{left}}), \mu(\mathsf{I}_j^{k,\mathsf{right}})) \ge \mu(L_{\ge j})/2$ (Claim 5.7). This proves part (i). To prove part (ii), let $x \prec y$ distinguish g_j from h, so $g_j(x) > g_j(y)$. We claim there exists a k^* such that $x \in \mathsf{I}_j^{k^*,\mathsf{left}}$ and $y \in \mathsf{I}_j^{k^*,\mathsf{right}}$. For any $\mathsf{I}_j^k = [a, b]$, the g_j values lie in $[2a+1, 2b+1]$. Hence, if $x \in \mathsf{I}_j^k$ and $y \notin \mathsf{I}_j^k$ (or vice versa), (x, y) is not a violation. So x and y lie in the same $\mathsf{I}_j^{k^*}$, But the function restricted to $\mathsf{I}_j^{k^*,\mathsf{left}}$ or $\mathsf{I}_j^{k^*,\mathsf{right}}$ is increasing, completing the proof. ☐

The following claim is a simple combinatorial statement about trees.

Claim 5.9. *Given a subset Q of $[n]$, let $\mathtt{lca}(Q) = \{\mathtt{lca}(x, y) : x, y \in Q\}$. Then $|\mathtt{lca}(Q)| \le |Q| - 1$.*

Proof. The proof is by induction on $|Q|$. The base case of $|Q| = 2$ is trivial. Suppose $|Q| > 2$. Consider the subset $P \subseteq Q$ of all elements of Q, none of whose ancestors are in Q. Also observe that if $P = Q$, then $\mathtt{lca}(Q)$ are precisely the internal nodes of a binary tree whose leaves are Q, and therefore $|\mathtt{lca}(Q)| \le |Q| - 1$. If P is a singleton, then $\mathtt{lca}(Q) = \mathtt{lca}(Q \setminus P) + 1 \le |Q \setminus P| - 1 + 1 = |Q| - 1$ (inequality from induction hypothesis). So assume $P \subset Q$ and $|P| \ne 1$. For $p \in P$, let S_p be the set of elements of Q appearing in the tree rooted at p. For every $x \in S_p$ and $y \in S_{p'}$ $(p \ne p')$, $\mathtt{lca}(x, y) = \mathtt{lca}(p, p')$. Furthermore, the sets S_p non-trivially partition Q. Therefore, $\mathtt{lca}(Q) = \mathtt{lca}(P) \cup \bigcup_{p \in P} \mathtt{lca}(S_p)$. Applying the induction hypothesis, $|\mathtt{lca}(Q)| \le |P| - 1 + \sum_{p \in P} |S_p| - |P| = |Q| - 1$. ☐

Let ℓ_ε be the largest ℓ such that $\mu(L_{\ge \ell}) \ge 2\varepsilon$. By Claim 5.8.(i), the collection of functions $\{g_1, \ldots, g_{\ell_\varepsilon}\}$ are each ε-far from monotone. By Claim 5.8.(ii) and Claim 5.9, a subset $Q \subseteq [n]$ can't distinguish more than $|Q|$ of these functions from h. Theorem 5.5 gives an $\Omega(\ell_\varepsilon)$ lower bound and Theorem 5.6 follows from Claim 5.10.

Claim 5.10. $\ell_\varepsilon \ge \rho\Delta^*(\mathscr{D})$.

Proof. Consider the distribution $\mathscr{D}'$ that transfers all the mass from $L_{\geq \ell_\varepsilon+1}$ to the remaining vertices proportionally. That is, if $\nu := \mu(L_{\geq \ell_\varepsilon+1})$, then $\mu_{\mathscr{D}'}(i) = 0$ for $i \in L_{\geq \ell_\varepsilon+1}$, and $\mu_{\mathscr{D}'}(i) = \mu_{\mathscr{D}}(i)/(1-\nu)$ for the rest. Observe that $||\mathscr{D} - \mathscr{D}'||_{\mathsf{TV}} = \mu_{\mathscr{D}}(L_{\geq \ell_\varepsilon+1}) < 2\varepsilon$. Also observe that since T is a binary tree of height ℓ_ε, $\ell_\varepsilon \geq \Delta^*(\mathscr{D}')$: the LHS is the max depth, the RHS is the (weighted) average depth. Now, we use stability of $\mathscr{D}$. Since $\mathscr{D}$ is $(2\varepsilon, \rho)$-stable, $\Delta^*(\mathscr{D}') \geq \rho\Delta^*(\mathscr{D})$. □

5.4 The Boolean hypercube

For the boolean hypercube, the lower bound doesn't require the stability assumption. Any product distribution over $\{0,1\}^d$ is determined by the d fractions $(\mu_1, \ldots, \mu_d)$, where μ_r is the probability of 0 on the r-th coordinate. Let $\theta_r := \min(\mu_r, 1-\mu_r)$.

Theorem 5.11. *Any monotonicity tester w.r.t. $\mathscr{D}$ for functions $f : \{0,1\}^d \mapsto \mathbb{N}$ with distance parameter $\varepsilon \leq 1/10$ must make $\Omega\left(\sum_{r=1}^d \min(\mu_r, 1-\mu_r)\right)$ queries.*

We begin with the basic setup. A tester for non-trivial ε makes at least 1 query, so we can assume that $\sum_{r=1}^d \theta_r > 1$. For ease of exposition, assume $\theta_r = \mu_r$, for all $1 \leq r \leq d$. (If not, we need to divide into two cases depending on θ_r and argue analogously for each case.) Assume wlog $\theta_1 \leq \theta_2 \leq \cdots \leq \theta_d$. Partition $[d]$ into contiguous segments $I_1, \ldots, I_b, I_{b+1}$ such that for each $1 \leq a \leq b$, $\sum_{r \in I_a} \theta_r \in [1/2, 1)$. Observe that $b = \Theta\left(\sum_r \theta_r\right)$. For $1 \leq a \leq b$, define the indicator functions $\chi_a : \{0,1\}^d \mapsto \{0,1\}$ as follows:

$$\chi_a(x) = \begin{cases} 1 & \text{if } \forall i \in I_a, x_i = 1 \\ 0 & \text{otherwise} \quad (\exists i \in I_a, x_i = 0) \end{cases}$$

By Theorem 5.5, we need to define the set of functions with appropriate properties. The monotone function $h(\cdot)$ is defined as $h(x) = \sum_{a=1}^b \chi_a(x) 2^a$. The functions $g_1, \ldots, g_b$ are defined as

$$g_a(x) = \begin{cases} h(x) - 2^r - 1 & \text{if } \chi_a(x) = 1 \\ h(x) & \text{if } \chi_a(x) = 0 \end{cases}$$

We prove all the desired properties.

Claim 5.12. *For all a, $\mathsf{dist}_{\mathscr{D}}(g_a, \mathtt{MON}) \geq 1/10$.*

Proof. Let I denote I_a, and $J = [n] \setminus I$. Think of $x = (x_I, x_J)$. Fix $\mathbf{v}$ in $\{0,1\}^{|J|}$, and define sets $X_1(\mathbf{v}) := \{x | \chi_a(x) = 1, x_J = \mathbf{v}\}$ (a singleton) and $X_0(\mathbf{v}) = \{x | \chi_a(x) = 0, x_J = \mathbf{v}\}$. Note that $\bigcup_{\mathbf{v}} (X_1(\mathbf{v}) \cup X_0(\mathbf{v}))$ forms a partition of the cube. For $c \neq a$, $\chi_c(x)$ is the same for all $x \in (X_0(\mathbf{v}) \cup X_1(\mathbf{v}))$. Hence, for *any* $x \in X_0(\mathbf{v})$ and $y \in X_1(\mathbf{v})$, $x \prec y$ and $g_a(x) > g_a(y)$. Any vertex cover in the violation graph must contain either $X_1(\mathbf{v})$ or $X_0(\mathbf{v})$, for each $\mathbf{v}$. Let $\mathscr{D}_I$ be the conditional distribution on the I-coordinates. In the following, we use the inequalities $\sum_{i \in I} \theta_i \in [1/2, 1)$ and $1 - t \in [e^{-2t}, e^t]$ for $t \leq 1/2$.

$$\begin{aligned} \mu_{\mathscr{D}_I}(X_1(\mathbf{v})) &= \prod_{i \in I}(1-\theta_i) \geq \exp(-2\sum_{i \in I}\theta_i) \geq e^{-2} > 1/10 \\ \mu_{\mathscr{D}_I}(X_0(\mathbf{v})) &= 1 - \prod_{i \in I}(1-\theta_i) \geq 1 - \exp(-\sum_{i \in I}\theta_i) \geq 1 - e^{-1/2} > 1/10. \end{aligned}$$

For each $\mathbf{v}$, the conditional mass of the vertex cover is at least $1/10$, and therefore, the $\mu_{\mathscr{D}}$ mass of the vertex cover is at least $1/10$. □

A pair x, y in $[n]^d$ *captures* index a if a is the largest index such that $\chi_a(x) \neq \chi_a(y)$. Furthermore, a set Q captures a if it contains a pair capturing a.

Claim 5.13. *If Q distinguishes g_a from h, then Q must capture a.*

Proof. Consider $x, y \in Q$ where $h(x) < h(y)$ but $g_a(x) > g_a(y)$. It must be that $\chi_a(x) = 0$ and $\chi_a(y) = 1$. Suppose this pair does not capture a. There must exist index $c > a$ (let it be the largest) such that $\chi_c(x) \neq \chi_c(y)$. Because $h(x) < h(y)$, $\chi_c(x) = 0$ and $\chi_c(y) = 1$. By definition, $g_a(y) - g_a(x) = (h(y) - 2^a - 1) - h(x)$. We have $h(y) - h(x) = \sum_{t=1}^{c}(\chi_t(y) - \chi_t(x))2^t$. Since $\chi_a(x) = \chi_c(x) = 0$ and $\chi_a(y) = \chi_c(y) = 1$, $h(y) - h(x) \geq 2^c + 2^a - \sum_{t<c:t\neq a} 2^t$. Combining,

$$g_a(y) - g_a(x) \geq 2^c + 2^a - \sum_{t<c:t\neq a} 2^t - 2^a - 1 = 2^a > 0.$$

□

Claim 5.14. *[Lifted from [CS13b].] A set Q captures at most $|Q| - 1$ coordinates.*

Proof. We prove this by induction on $|Q|$. When $|Q| = 2$, this is trivially true. Otherwise, pick the largest coordinate j captured by Q and let $Q_0 = \{x : x_j = 0\}$ and $Q_1 = \{x : x_j = 1\}$. By induction, Q_0 captures at most $|Q_0| - 1$ coordinates, and Q_1 captures at most $|Q_1| - 1$ coordinates. Pairs $(x, y) \in Q_0 \times Q_1$ only capture coordinate j. The total number of captured coordinates is at most $|Q_0| - 1 + |Q_1| - 1 + 1 = |Q| - 1$. □

We can now invoke Theorem 5.5 to get an $\Omega(b) = \Omega(\sum_r \theta_r)$ lower bound thereby proving Theorem 5.11. The hypercube lower bound can be generalized to give a weak lower bound for hypergrids, which will be useful for proving the stronger bound. Fix a dimension r. For any $1 \leq j \leq n$, define $\theta_r^j := \min(\sum_{k\leq j} \mu_{\mathscr{D}_r}(k), 1 - \sum_{k\leq j} \mu_{\mathscr{D}_r}(k))$. Define $\theta_r := \mathsf{max}_{1\leq j\leq n}\theta_r^j$. Note that θ_r generalizes the above definition for the hypercube. The following theorem follows by a reduction to the hypercube lower bound.

Theorem 5.15. *Any monotonicity tester on the hypergrid with distance parameter $\varepsilon \leq 1/10$, makes $\Omega\left(\sum_{r=1}^{d} \theta_r\right)$ queries.*

Proof. For $1 \leq r \leq d$, let $1 \leq j_r \leq n$ be the j such that $\theta_r = \theta_r^j$. Project the hypergrid onto a Boolean hypercube using the following mapping $\psi : [n]^d \to \{0,1\}^d$: for $x \in [n]^d$, $\psi(x)_r = 0$ if $x_r \leq j_r$, and 1 otherwise. The corresponding product distribution $\mathscr{D}'$ on the hypercube puts $\mu_{\mathscr{D}'_r}(0) = \sum_{k\leq j_r} \mu_{\mathscr{D}_r}(k)$, for all r. Note that $\min(\mu_r, 1 - \mu_r) = \theta_r$. Given any function f on $\{0,1\}^d$, extend it to g over the hypergrid in the natural way: for $x \in [n]^d$, $g(x) = f(\psi(x))$. Note that $\mathsf{dist}_{\mathscr{D}'}(f, \mathtt{MON}) = \mathsf{dist}_{\mathscr{D}}(g, \mathtt{MON})$. (This is akin to Theorem 3.8.) Any tester for g over $[n]^d$ induces a tester for f on $\{0,1\}^d$ with as good a query complexity: whenever the hypergrid tester queries $x \in [n]^d$, the hypercube tester queries $\psi(x)$. Therefore, the lower bound Theorem 5.11 for the hypercube implies Theorem 5.15. □

5.5 The Hypergrid.

Our main lower bound result is the following, which implies Theorem 1.6 via Theorem 5.4.

Theorem 5.16. *For any parameter $\varepsilon < 1/10$, and for any $(120\varepsilon, \rho)$-stable, product distribution $\mathscr{D}$, any (even adaptive, two-sided) montonicity tester w.r.t. $\mathscr{D}$ for functions $f : [n]^d \mapsto \mathbb{N}$ with proximity parameter ε requires $\Omega(\rho\Delta^*(\mathscr{D}))$ queries.*

5.5.1 The intuition

Since we already have a proof for $d = 1$ in §5.3, an obvious approach to prove Theorem 1.6 is via some form of induction on the dimension. Any of the g_j-functions on $[n]$ in §5.3 can be extended the obvious way to a function on $[n]^d$. Given (say) $g_j : [n] \mapsto \mathbb{N}$, we can define $f : [n]^d \mapsto \mathbb{N}$ as $f(x) = g_j(x_1)$. Thus, we embed the hard functions for $\mathscr{D}_1$ along dimension 1. One can envisage a way do the same for dimension 2, and so on and so forth, thereby leading to $\sum_i \Delta^*(\mathscr{D}_i)$ hard functions in all. There is a caveat here. The construction of §5.3 for (say) $\mathscr{D}_1$ requires the stability of $\mathscr{D}_1$. Otherwise, we don't necessarily get $\Omega(\Delta^*(\mathscr{D}_1))$ functions with distance at least ε. For instance, if the root of the median BST has more than $(1-\varepsilon)$ fraction of the weight, we get at most one hard function of distance at least ε. So, the above approach requires stability of all the *marginals* of $\mathscr{D}$. Unfortunately, there exist stable product distributions with all marginals unstable. Consider $\mathscr{D} = \prod_r \mathscr{D}_r$, where each $\mathscr{D}_r = (\frac{1}{(n-1)d}, \ldots, \frac{1}{(n-1)d}, 1 - \frac{1}{d})$. Note that $\Delta^*(\mathscr{D}) \approx \log n$. Each $\mathscr{D}_r$ is individually unstable (for $\varepsilon > 1/d$), since there is a $\mathscr{D}'_i$ with all the mass on the nth coordinate, such that $\|\mathscr{D}_i - \mathscr{D}'_i\|_{TV} = 1/d$ and $\Delta^*(\mathscr{D}'_i) = 0$. On the other hand, it is not hard to see that $\mathscr{D}$ is $(1/100, 1/100)$-stable. A new idea is required to construct the lower bound. To see this, suppose there is a product distribution $\mathscr{D}'$ such that $\Delta^*(\mathscr{D}') < \Delta^*(\mathscr{D})/100 = (\log n)/100$. Markov's inequality implies that for $\Omega(d)$ dimensions, $\|\mathscr{D}'_r - \mathscr{D}_r\|_{\mathsf{TV}} = \Omega(1/d)$. A calculation shows that $\|\mathscr{D} - \mathscr{D}'\|_{\mathsf{TV}}$ must be at least $1/100$. In sum, for any constants ε, ρ, there exist (ε, ρ)-stable distributions $\mathscr{D}$ such that each marginal $\mathscr{D}_r$ is only $(\varepsilon/d, \rho)$-stable. This is a major roadblock for a lower bound construction, and therefore a new idea is required. We design an *aggregation* technique that does the following. Start with 1D functions $g^1_{j_1}$ and $g^2_{j_2}$ that are hard functions from §5.3 for $\mathscr{D}_1$ and $\mathscr{D}_2$ respectively. Suppose the corresponding distances to monotonicity are $\varepsilon^{(1)}$ and $\varepsilon^{(2)}$. We construct a function $f : [n]^d \mapsto \mathbb{N}$ that is $\varepsilon^{(1)} + \varepsilon^{(2)}$-far, so we can effectively add their distances. If we can aggregate $\Omega(d)$ 1D functions, each with distance ε/d, then we get a desired hard function. As can be expected, this construction is quite delicate, because we embed violations in many dimensions simultaneouly. Furthermore, we need to argue that this aggregation can produce enough "independent" hard functions, so we get a large enough lower bound (from Theorem 5.5). And that is where the hard work lies.

5.5.2 Setup and Construction

Fix ε and let $\varepsilon' = 120\varepsilon$. Fix the (ε', ρ)-stable distribution $\mathscr{D}$. Since $\mathscr{D}$ is (ε', ρ)-stable, for any $\mathscr{D}'$ with $\|\mathscr{D}' - \mathscr{D}\|_{\mathsf{TV}} \le \varepsilon'$, we have $\Delta^*(\mathscr{D}') \ge \rho\Delta^*(\mathscr{D})$. We denote the median BST for $\mathscr{D}_r$ as T_r, Δ_r as the expected depth w.r.t. $\mathscr{D}_r$, and $\Delta(\mathscr{D}) = \sum_{r=1}^d \Delta_r$. The following shows that the median BST is near optimal.

Lemma 5.17. *For any product distribution $\mathscr{D} = \prod_r \mathscr{D}_r$, $\Delta(\mathscr{D}) \le 5\Delta^*(\mathscr{D})$.*

Proof. Fix a coordinate r. The depth of a vertex u in T_r is at most $\log_2(1/\mu_{\mathscr{D}_r}(u))$, so we get $\Delta_r \le H(\mathscr{D}_r)$, the Shannon entropy of $\mathscr{D}_r$. It is also known (cf. Thm 2 in [Meh75]) that $H(\mathscr{D}_r) \le \log_2 3(\Delta^*(\mathscr{D}_r) + 1)$. To see this, notice that any BST can be converted into a prefix-free ternary code of expected length $(\Delta^*(\mathscr{D}_r) + 1)$, say, over the alphabet 'left','right', and 'stop'. Therefore, if $\Delta^*(\mathscr{D}_r) \ge 1/2$, we have $\Delta_r \le 5\Delta^*(\mathscr{D}_r)$. If $\Delta^*(\mathscr{D}_r) < 1/2$, then since $\Delta^*(\mathscr{D}_r) \ge 1 - \Pr[\text{root}]$, we get $\mu^* := \mu_{\mathscr{D}_r}(u^*) > 1/2$ where u^* is the root of the optimal BST T^*. But this implies u^* is also the root of T_r by construction of the median BST. Now we can prove via induction. If p and q are the total masses of the nodes in the left and right sub-tree of T^* (and therefore also T_r), and Δ^*_1

(resp. Δ_1) and Δ_2^* (resp. Δ_2)be the expected depths of these subtrees in T^* (resp. T_r), then we get, $\Delta^*(\mathscr{D}_r) = p\Delta_1^* + q\Delta_2^* + (1-\mu^*) \leq 5p\Delta_1 + 5q\Delta_2 + (1-\mu^*) \leq 5\Delta_r$. □

Theorem 5.5 requires the definition of a monotone function and a collection of ε-far from monotone functions with additional properties. The monotone function is $\mathsf{val}(x) := \sum_{r=1}^{d} 2(2n+1)^r x_r$. The non-monotone functions (which we refer to as "hard" functions) are constructed via aggregation. From §5.3, for each dimension r and each level $j \geq 1$ in tree T_r, we have a 1D "hard" function $g_j^r : [n] \mapsto \mathbb{N}$. It is useful to abstract out some of the properties of g_j^r that were proved in §5.3. Let L_j^r be the nodes in T_r at level j. Each level corresponds to a collection of intervals of $[n]$. We use $L_{\geq j}^r := \bigcup_{j' \geq j} L_{j'}^r$ and $L_{<j}^r = \bigcup_{j'<j} L_{j'}^r$. We use the shorthand $\mu_{\geq j}^r$ to denote $\mu_{\mathscr{D}_r}(L_{\geq j}^r)$. The following lemma is a restatement of Claim 5.7 and Claim 5.8.

Lemma 5.18. *Consider $g_j^r : [n] \mapsto \mathbb{N}$, for $j \geq 1$ All violations to monotonicity are contained in intervals corresponding to L_{j-1}^r, and the distance to monotonicity is at least $\mu_{\geq j}^r/2$. Furthermore, any violation (x, y) has $\mathtt{lca}(x, y)$ in L_{j-1}^r.*

The aggregation process takes as input a *map* $\psi : [d] \mapsto \{\perp\} \cup \{2, 3, 4, \ldots\}$. Note that if $\psi(r) \neq \perp$, then $\psi(r) > 1$. Informally, $\psi(r)$, when not equating to $\perp$, tells us the level of T_r whose hard function is to be included in the aggregation. We define $\Psi^{-1} := \{r | \psi(r) \neq \perp\}$, the subset of relevant dimensions. Given the map ψ, we aggregate the collection of 1D functions $\{g_{\psi(r)}^r | r \in \Psi^{-1}\}$ into a single hard function for $[n]^d$ as follows.

$$g_\psi(x) := \sum_{r \in \Psi^{-1}} (2n+1)^r g_{\psi(r)}^r(x_r) + \sum_{r \notin \Psi^{-1}} 2(2n+1)^r x_r \tag{15}$$

Observe that the latter sum is identical to the corresponding portion in $\mathsf{val}(x)$. The first summand takes the hard function corresponding to the $\psi(r)$th level of T_r for $r \in \Psi^{-1}$ and aggregates them via multiplying them with a suitable power of $(2n+1)$.

Definition 5.19. *A map ψ is* **useful** *if the following are true.*

- $\sum_{r \in \Psi^{-1}} \mu_{\geq \psi(r)}^r \in (\varepsilon', 1)$
- *For all $r \in \Psi^{-1}$, $\mu_{\geq \psi(r)}^r \geq \frac{\mu_{\geq \psi(r)-1}^r}{2}$.*

In plain English, the first point states that total distance of the hard functions picked should be at least ε'. The second point is a technicality which is required to argue about the distance of the aggregated function. It states that in each relevant T_r, the total mass on the nodes lying in the $\psi(r)$th layer and below shouldn't be much smaller than the total mass on the nodes lying on the $(\psi(r)-1)$th layer and below.

Lemma 5.20. *If ψ is useful, $\mathsf{dist}_{\mathscr{D}}(g_\psi, \mathtt{MON}) \geq \varepsilon$.*

Proof. It is convenient to consider restrictions of g_ψ where all coordinates in $[d] \setminus \Psi^{-1}$ are fixed. This gives rise to $|\Psi^{-1}|$-dimensional functions. We argue that each such restriction is ε-far from monotone, which proves the lemma. Abusing notation, we use g_ψ to refer to an arbitrary such restriction. Fix some $r \in \Psi^{-1}$. Define the subset $S_r := \{x \in [n]^d : x_s \in L_{<\psi(s)-1}^s, \forall s \neq r\}$ to be the set of points x with the sth coordinate appearing in the first $(\psi(s)-2)$ layers of the tree T_s, for all $s \neq r$. We stress that this is well-defined because $\psi(s) \geq 2$ by definition of ψ. Note that each S_r is a collection of r-lines and the restriction of g_ψ on each line exactly a multiple of $g_{\psi(r)}^r$. By Lemma 5.18,

all violations to monotonicity in such lines lie in the intervals corresponding to $L^r_{\geq\psi(r)-1}$, and the mass of the vertex cover of the violation graph (restricted to the line) is at least $\mu^r_{\geq\psi(r)}/2$. Thus the total contribution to distance of g_ψ from S_r is at least $\frac{\mu^r_{\geq\psi(r)}}{2} \cdot \mu_{\mathscr{D}_{-r}}(\prod_{s\neq r} L^s_{<\psi(s)-1})$. What is crucial to note is that the regions of violations in S_r is *disjoint* from the regions of violation in $S_{r'}$ for $r' \neq r$. Therefore, the contributions to the distance of g_ψ add up, and this gives

$$\begin{aligned}
\mathsf{dist}_{\mathscr{D}}(g_\psi, \texttt{MON}) &\geq \frac{1}{2}\sum_{r\in\Psi^{-1}} \mu^r_{\geq\psi(r)} \cdot \mu_{\mathscr{D}_{-r}}\Big(\prod_{s\neq r} L^s_{<\psi(s)-1}\Big) \\
&= \frac{1}{2}\sum_{r\in\Psi^{-1}} \mu^r_{\geq\psi(r)} \prod_{s\neq r}(1-\mu^s_{\geq\psi(s)-1}) \\
&\geq \frac{1}{2}\sum_{r\in\Psi^{-1}} \mu^r_{\geq\psi(r)} \prod_{s\neq r}(1-2\mu^s_{\geq\psi(s)}) \qquad \text{(point 2 in def. of useful map)}
\end{aligned}$$

We can apply the bound, $\sum_{r\in\Psi^{-1}} \mu^r_{\geq\psi(r)} \in (\varepsilon', 1)$, since ψ is useful. We lower bound the product by $\exp(-4\sum_{s\neq r}\mu^s_{\geq\psi(s)})$, which by the above bound, is at least e^{-4}. So, $\mathsf{dist}_{\mathscr{D}}(g_\psi, \texttt{MON}) \geq \sum_{r\in\Psi^{-1}} \mu^r_{\geq\psi(r)}/120 \geq \varepsilon'/120 = \varepsilon$. □

Definition 5.21. *Two maps ψ_1, ψ_2 are* disjoint *if: $\{(r, \psi_1(r)) | r \in \Psi_1^{-1}\}$ and $\{(r, \psi_2(r)) | r \in \Psi_2^{-1}\}$ are disjoint.*

That is, for every tree T_r, ψ_1 and ψ_2 point to different layers of the tree (or they point to $\perp$).

Lemma 5.22. *Consider a set of maps $\psi_1, \psi_2, \ldots$ that are all pairwise disjoint. A set of Q queries can distinguish at most $|Q| - 1$ of these functions from* val.

Proof. Say a pair (x, y) of queries captures the (unique) tuple (r, j) if the largest coordinate in which x and y differ is r, and furthermore $\texttt{lca}(x_r, y_r)$ in T_r lies in level $(j-1)$. A set Q captures (r, j) if some pair in Q captures (r, j). We first show that if (x, y) distinguishes g_ψ from val for some map ψ, then (x, y) captures a pair $(r, \psi(r))$ for some $r \in \Psi^{-1}$. Assume wlog $\mathsf{val}(x) < \mathsf{val}(y)$, and so $g_\psi(x) > g_\psi(y)$. Let a be the largest coordinate at which x and y differ; since $\mathsf{val}(x) < \mathsf{val}(y)$, we get $x_a < y_a$. Suppose $g^a_{\psi(a)}(x_a)$ and $g^a_{\psi(a)}(y_a)$ is not a violation. By the construction, this implies that $g^a_{\psi(a)}(y_a) - g^a_{\psi(a)}(x_a) \geq 1$. Furthermore, $g^r_{\psi(r)}$ is always in the range $[1, 2n]$ for any r.

$$\begin{aligned}
g_\psi(y) - g_\psi(x) &= (2n+1)^a(g^a_{\psi(a)}(y_a) - g^a_{\psi(a)}(x_a)) + \sum_{r<a, r\in\Psi^{-1}} (2n+1)^r(g^r_{\psi(r)}(y_r) - g^r_{\psi(r)}(x_r)) \\
&\geq (2n+1)^a - (2n)\sum_{r<a}(2n+1)^r \\
&= (2n+1)^a - (2n)\cdot\frac{(2n+1)^a - 1}{2n} \qquad > 0,
\end{aligned}$$

So $(g^a_{\psi(a)}(x_a), g^a_{\psi(a)}(y_a))$ is a violation. Immediately, we deduce that $\psi(a) \neq \perp$, so $a \in \Psi^{-1}$. By Claim 5.8, $\texttt{lca}(x_a, y_a)$ lies in level $\psi(a) - 1$ of T_a, and hence, (x, y) captures $(a, \psi(a))$. As we prove in Claim 5.23, Q queries can capture at most $|Q| - 1$ such tuples. The proof is completed by noting the maps $\psi_1, \psi_2, \ldots$ are pairwise disjoint. □

Claim 5.23. *A nonempty set Q can only capture at most $|Q| - 1$ tuples (r, j).*

Proof. Proof is by induction on $|Q|$. If $|Q| = 2$, then the claim trivially holds. Assume $|Q| > 2$. Let s be the largest dimension such that there are at least two points in Q differing in that dimension. For $c = 1$ to n, let $Q_c := \{x \in Q : x_s = c\}$. By definition, $Q_c \subset Q$. Reorder the dimensions such that Q_c is non-empty for $c = 1 \dots q \leq n$. By induction, each Q_c captures at most $|Q_c| - 1$ pairs for $1 \leq c \leq q$. Consider (x, y) with $x \in Q_c$ and $y \in Q_{c'}$ for $c \neq c'$. The largest coordinate where they differ is exactly s. All tuples captured by such pairs is of the form (s, ℓ), where ℓ is the `lca` in T_s of some $c, c' \in \{1 \dots, q\}$. By Claim 5.9, the total number of such points is at most $q - 1$. Thus, the total number of tuples captured is at most $\sum_{a=1}^{q} |Q_a| - q + (q - 1) = |Q| - 1$. □

5.5.3 Constructing the maps

Let us go back to the framework of Theorem 5.5. From Lemma 5.20 and Lemma 5.22, it suffices to construct a sequence $\psi_1, \psi_2, \dots$ of pairwise disjoint, useful maps. The number of such maps will exactly be our lower bound. The exact construction is a little tricky, since the conditions of usefulness are somewhat cumbersone. We use the following definition.

- **Allowed levels:** A level j is allowed w.r.t. dimension r if $j > 1$ and $\mu^r_{\geq j} \geq \mu^r_{\geq j-1}/2$. This is in lines with point 2 of the usefulness definition.
- **Level sets A_r:** A_r is the set of allowed levels of tree T_r.

It is convenient to define an abstract procedure that constructs these maps. We have a stack S_r for each $r \in [d]$, whose elements are allowed levels. The stack S_r is initialized with A_r in increasing order, that is the head (top entry) of the stack is the least (that is, closest to root) level in A_r. In each *round*, we will construct a map ψ. Denote the head of S_r by h_r. Note that $h_r > 1$ by definition of allowed levels. Maintain a running count initialized to 0. We go through the stacks in an arbitrary order popping off a *single* element from each stack. In a round, we never touch the same stack more than once. When we pop S_r, we set $\psi(r) := h_r$ and add $\mu^r_{\geq h_r}$ to the running count. We stop as soon as the running count enters the interval $[\varepsilon', 1]$. For all r for which $\psi(r)$ hasn't been defined, we set $\psi(r) = \perp$. This completes the description of a single map. Observe, by definition of allowed levels and the stopping condition, ψ is useful. When $\sum_{r=1}^{d} \mu^r_{\geq h_r} < \varepsilon'$, we cannot complete the construction. So the procedure terminates, discarding the final map. Let the set of maps constructed be $\mathbf{\Psi}$. By construction, the maps are useful. Furthermore, they are pairwise disjoint, because once a layer is popped out, it never appears again. We now basically show that $|\mathbf{\Psi}|$ is large, using the (ε', ρ)-stability of $\mathscr{D}$. This proves that the number of hard functions is large. We have to first deal with an annoying corner case of $\mathscr{D}$.

Theorem 5.24. *If $\sum_r \mu^r_{\geq 1} > \rho\Delta(\mathscr{D})/12$, then any monotonicity tester requires $\Omega(\rho\Delta^*(\mathscr{D}))$ queries.*

Proof. We simply apply the hypercube lower bound. Recall the definition of θ_r described before Theorem 5.15. Note that $\mu^r_{\geq 1}$ is simply the total $\mathscr{D}_r$-mass of everything in T_r other than the root. By the median property of the T_r, θ_r is ensured to be at least half of this mass, and hence $\theta_r \geq \mu^r_{\geq 1}/2$. Combining with Theorem 5.15, we get a lower bound of $\Omega(\sum_r \mu^r_{\geq 1})$, which by assumption, is $\Omega(\rho\Delta(\mathscr{D}))$. An application of Lemma 5.17 completes the proof. □

Now we come to the main bound of $|\mathbf{\Psi}|$. We need some setup for the proof. The following simple observation is crucial. This follows since $\mathbf{E}[Z] = \sum_{k \in \mathbb{N}} \Pr[Z \geq k]$, for any non-negative, integer valued random variable.

Claim 5.25. *For all r, $\sum_{j \geq 1} \mu^r_{\geq j} = \mathbf{E}_{x \sim \mathscr{D}_r}[\mathsf{depth}_{T_r}(x)] = \Delta_r$.*

The following lemma completes the entire lower bound.

Lemma 5.26. *Suppose* $\sum_r \mu^r_{\geq 1} \leq \rho\Delta(\mathscr{D})/12$. *Then* $|\mathbf{\Psi}| = \Omega(\rho\Delta(\mathscr{D}))$.

Proof. Let h_r denote the head of S_r when the procedure terminates. So, $\sum_{r=1}^d \mu^r_{\geq h_r} < \varepsilon'$. For any $\psi \in \mathbf{\Psi}$, $\sum_{r \in \Psi^{-1}} \mu^r_{\geq \psi(r)} \leq 1$. Hence, $|\mathbf{\Psi}|$ is at least the total sum over popped elements $\mu^r_{\geq j}$. Writing this out and expanding out a summation,

$$|\mathbf{\Psi}| \geq \sum_{r\in[d]} \sum_{j<h_r, j\in A_r} \mu^r_{\geq j} \quad = \quad \sum_{r\in[d]} \Big[\sum_{j=1}^{h_r-1} \mu^r_{\geq j} - \mu^r_{\geq 1} - \sum_{1<j<h_r : j\notin A_r} \mu^r_{\geq j} \Big]$$

Recall that $h_r > 1$ and so the summations are well-defined. For any level $1 < j \notin A_r$, we have $\mu^r_{\geq j} < \mu^r_{\geq j-1}/2$. Therefore, $\sum_{1<j<h_r : j\notin A_r} \mu^r_{\geq j} < \sum_{j=1}^{h_r-2} \mu^r_{\geq j}/2$. Plugging this bound in and applying the lemma assumption,

$$|\mathbf{\Psi}| \geq \sum_{r\in[d]} \sum_{j=1}^{h_r-1} \mu^r_{\geq j}/2 - \sum_{r\in[d]} \mu^r_{\geq 1} \geq \sum_{r\in[d]} \sum_{j=1}^{h_r-1} \mu^r_{\geq j}/2 - \rho\Delta(\mathscr{D})/12 \tag{16}$$

We need to lower bound the double summation above. Observe that the second summation is $\sum_{j\geq 1} \mu^r_{\geq j} - (\mu^r_{\geq h_r} + \mu^r_{\geq h_r+1} + \cdots)$. The first term, by Claim 5.25 is precisely $\Delta(\mathscr{D})$, and by definition of h_r, each of the terms in the parenthesis is at most ε'. However, the number of terms in the parenthesis can be quite large, and this doesn't seem to get any lower bound on the summation. Here's where stability of $\mathscr{D}$ saves the day. Construct a distribution $\mathscr{D}'_r$ be the distribution on $[n]$ as follows. Move the entire probability mass away from $L^r_{\geq h_r}$ and distribute it on the ancestral nodes in level $(h_r - 1)$ of T_r. More precisely, $\mu_{\mathscr{D}'_r}(u) = 0$ if $u \in L^r_{\geq h_r}$, $\mu_{\mathscr{D}'_r}(u) = \mu_{\mathscr{D}_r}(u)$ if $u \in L^r_{<h_r-1}$, and $\mu_{\mathscr{D}'_r}(u) = \sum_v \mu_{\mathscr{D}_r}(v)$ for $u \in L^r_{h_r-1}$ where the summation is over children v of u in T_r. Letting $\mathscr{D}' := \prod_r \mathscr{D}'_i$, we see that $||\mathscr{D} - \mathscr{D}'||_{\mathsf{TV}} \leq \sum_{r=1}^d \mu^r_{\geq h_r} < \varepsilon'$. Since $\mathscr{D}$ is (ε', ρ)-stable, we get $\Delta^*(\mathscr{D}') \geq \rho\Delta^*(\mathscr{D})$. Now we can apply Claim 5.25 on $\mathscr{D}'_r$ and T_r to get $\mathbf{E}_{x\sim\mathscr{D}'_r}[\mathsf{depth}_{T_r}(x)] = \sum_{j\geq 1} \mu_{\mathscr{D}'_r}(L^r_{\geq j}) = \sum_{j=1}^{h_r-1} \mu^r_{\geq j}$. This expected depth is by definition at least $\Delta^*(\mathscr{D}'_r)$. Therefore, we get a lower bound of $\sum_{r=1}^d \Delta^*(\mathscr{D}'_r)/2 = \Delta^*(\mathscr{D}')/2$ on the double summation in (16). Using the stability of $\mathscr{D}'$ this is at least $\rho\Delta^*(\mathscr{D})/2$. Substituting we get

$$\begin{aligned} |\mathbf{\Psi}| &\geq \rho\Delta^*(\mathscr{D})/2 - \rho\Delta(\mathscr{D})/12 \\ &\geq \rho\Delta^*(\mathscr{D})/2 - 5\rho\Delta^*(\mathscr{D})/12 = \Omega(\rho\Delta^*(\mathscr{D})) \qquad \text{(by Lemma 5.17)} \end{aligned}$$

□

We put it all together to prove the main lower bound, Theorem 5.16. If $\sum_r \mu^r_{\geq 1} > \rho\Delta(\mathscr{D})/12$, Theorem 5.24 proves Theorem 5.16. Otherwise, by Lemma 5.26 we have constructed $\Omega(\rho\Delta^*(\mathscr{D}))$ pairwise disjoint, useful maps. Each map yields a hard function of distance at least ε (by Lemma 5.20), and these functions satisfy the conditions of Theorem 5.5, which implies Theorem 5.16.

References

[AC06] N. Ailon and B. Chazelle. Information theory in property testing and monotonicity testing in higher dimension. *Inform. and Comput.*, 204(11):1704–1717, 2006. 2, 3, 4, 5, 6

[ACCL07] N. Ailon, B. Chazelle, S. Comandur, and D. Liu. Estimating the distance to a monotone function. *Random Structures Algorithms*, 31(3):371–383, 2007. 2, 5, 15

[AFNS09] N. Alon, E. Fischer, I. Newman, and A. Shapira. A combinatorial characterization of the testable graph properties : it's all about regularity. *SIAM J. Comput.*, 39(1):143–167, 2009. 1

[AJMR12] P. Awasthi, M. Jha, M. Molinaro, and S. Raskhodnikova. Testing Lipschitz functions on hypergrid domains. In *Proceedings, International Workshop on Randomization and Computation (RANDOM)*, 2012. 2, 4

[AS08] N. Alon and A. Shapira. A charaterization of the (natural) graph properties testable with one-sided error. *SIAM J. Comput.*, 37(6):1703–1727, 2008. 1

[BBBY12] M-F. Balcan, E. Blais, A. Blum, and L. Yang. Active property testing. In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, 2012. 6

[BBM12] E. Blais, J. Brody, and K. Matulef. Property testing lower bounds via communication complexity. *Comp. Complexity*, 21(2):311–358, 2012. 2, 3

[BCGSM12] J. Briët, S. Chakraborty, D. García-Soriano, and A. Matsliah. Monotonicity testing and shortest-path routing on the cube. *Combinatorica*, 32(1):35–53, 2012. 2

[BGJ+09] A. Bhattacharyya, E. Grigorescu, K. Jung, S. Raskhodnikova, and D. Woodruff. Transitive-closure spanners. In *Proceedings, ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 531–540, 2009. 2, 5, 15

[BRW05] T. Batu, R. Rubinfeld, and P. White. Fast approximate PCPs for multidimensional bin-packing problems. *Inform. and Comput.*, 196(1):42–56, 2005. 2, 5, 15

[BRY14a] P. Berman, S. Raskhodnikova, and G. Yaroslavtsev. Testing with respect to l_p distances. In *Proceedings, ACM Symp. on Theory of Computing (STOC)*, 2014. 6

[BRY14b] E. Blais, S. Raskhodnikova, and G. Yaroslavtsev. Lower bounds for testing properties of functions on hypergrid domains. In *Proceedings, IEEE Conference on Computational Complexity (CCC)*, March 2014. 2, 4, 5

[BSS10] I. Benjamini, O. Schramm, and A. Shapira. Every minor-closed property of sparse graphs is testable. *Adv. in Math.*, 223(6):2200–2218, 2010. 1

[CS13a] D. Chakrabarty and C. Seshadhri. Optimal bounds for monotonicity and Lipschitz testing over hypercubes and hypergrids. In *Proceedings, ACM Symp. on Theory of Computing (STOC)*, 2013. 2, 3, 5, 6, 11

[CS13b] D. Chakrabarty and C. Seshadhri. An optimal lower bound for monotonicity testing over hypergrids. In *Proceedings, International Workshop on Randomization and Computation (RANDOM)*, 2013. 2, 3, 5, 6, 19, 22

[DGL+99] Y. Dodis, O. Goldreich, E. Lehman, S. Raskhodnikova, D. Ron, and A. Samorodnitsky. Improved testing algorithms for monotonicity. In *Proceedings, International Workshop on Randomization and Computation (RANDOM)*, 1999. 2, 3, 4

[DJRT13] K. Dixit, M. Jha, S. Raskhodnikova, and A.G. Thakurta. Testing the Lipschitz property over product distributions with applications to data privacy. In *Proceedings, Theory of Cryptography Conference (TCC)*, 2013. 1, 2

[DR11] E. Dolev and D. Ron. Distribution-free testing for monomials with a sublinear number of queries. *Theory of Computing*, 7(1):155–176, 2011. 6

[EKK+00] F. Ergun, S. Kannan, R. Kumar, R. Rubinfeld, and M. Viswanathan. Spot-checkers. *J. Comput. System Sci.*, 60(3):717–751, 2000. 2, 5, 15

[Fis04] E. Fischer. On the strength of comparisons in property testing. *Inform. and Comput.*, 189(1):107–116, 2004. 2, 5, 19

[FLN+02] E. Fischer, E. Lehman, I. Newman, S. Raskhodnikova, R. Rubinfeld, and A. Samorodnitsky. Monotonicity testing over general poset domains. In *Proceedings, ACM Symp. on Theory of Computing (STOC)*, 2002. 2, 3, 8

[GGL+00] O. Goldreich, S. Goldwasser, E. Lehman, D. Ron, and A. Samorodnitsky. Testing monotonicity. *Combinatorica*, 20:301–337, 2000. 2, 3, 4

[GGR98] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *J. ACM*, 45(4):653–750, 1998. 1, 6

[GR13] O. Goldreich and D. Ron. On sample-based testers. *Elect. Coll. on Comp. Complexity (ECCC)*, 20:109, 2013. 6

[GS09] D. Glasner and R. A. Servedio. Distribution-free testing lower bound for basic boolean functions. *Theory of Computing*, 5(1):191–216, 2009. 6

[Hal06] S. Halevy. *Topics in Property Testing*. PhD thesis, Tel Aviv University, 2006. 6

[HK05] S. Halevy and E. Kushilevitz. A lower bound for distribution-free monotonicity testing. In *Proceedings, International Workshop on Randomization and Computation (RANDOM)*, pages 330–341, 2005. 6

[HK07] S. Halevy and E. Kushilevitz. Distribution-free property-testing. *SIAM J. Comput.*, 37(4):1107–1138, 2007. 2, 6

[HK08a] S. Halevy and E. Kushilevitz. Distribution-free connectivity testing for sparse graphs. *Algorithmica*, 51(1):24–48, 2008. 6

[HK08b] S. Halevy and E. Kushilevitz. Testing monotonicity over graph products. *Random Structures Algorithms*, 33(1):44–67, 2008. 2, 4, 5, 6

[JR11] M. Jha and S. Raskhodnikova. Testing and reconstruction of Lipschitz functions with applications to data privacy. In *Proceedings, IEEE Symposium on Foundations of Computer Science (FOCS)*, 2011. 2, 4, 5, 15

[Knu73] D. E. Knuth. *The Art of Computer Programming Vol III: Sorting and Searching*, volume 3. Addison-Wesley, 1973. 3

[KS08a] T. Kaufman and M. Sudan. Algebraic property testing: the role of invariance. In *Proceedings, ACM Symp. on Theory of Computing (STOC)*, 2008. 1

[KS08b] S. Kopparty and S. Saraf. Tolerant linearity testing and locally testable codes. In *Proceedings, International Workshop on Randomization and Computation (RANDOM)*, 2008. 6

[LR01] E. Lehman and D. Ron. On disjoint chains of subsets. *J. Combin. Theory Ser. A*, 94(2):399–404, 2001. 2, 3

[Meh75] K. Mehlhorn. Nearly optimal binary search trees. *Acta Informatica*, 5:287–295, 1975. 3, 19, 23

[PRR06] M. Parnas, D. Ron, and R. Rubinfeld. Tolerant property testing and distance approximation. *J. Comput. System Sci.*, 6(72):1012–1042, 2006. 2, 5

[ver14] Random constraints in systemverilog. http://www.asic-world.com/systemverilog/random_const 2014. 1

[VHD02] Systemverilog 3.1: Random constraints - proposal. http://www.vhdl.org/sv-ec/SV_3.1_Web/Random-Constraints_Proposal.pdf, 2002. 1

[Yao82] F. F. Yao. Speed-up in dynamic programming. *J. Alg. Discrete Math.*, 3:532–540, 1982. 3